arXiv:1701.02643v1 [stat.AP] 10 Jan 2017

Bayesian log-Gaussian Cox process regression with applications to fMRI meta-analysis

Pantelis Samartsidis[1], Tor D. Wager[2], Lisa Feldman Barrett[3], Shir Atzil[4], Simon B. Eickhoff[5], Thomas E. Nichols[6] and Timothy D. Johnson[7]

[1]*MRC Biostatistics Unit* [2]*University of Colorado at Boulder*
[3]*Northeastern University* [4]*Massachusetts General Hospital*
[5]*Heinrich-Heine University Düsseldorf* [6]*University of Warwick*
[7]*University of Michigan*

Abstract

A typical neuroimaging study will produce a 3D brain statistic image that summarises the evidence for activation during the experiment. However, for practical reasons those images are rarely published; instead, authors only report the (x,y,z) locations of local maxima in the statistic image. Neuroimaging meta-analyses use these foci from multiple studies to find areas of consistent activation across the human brain. However, current methods in the field do not account for study-specific characteristics. Therefore, we propose a fully Bayesian model based on log-Gaussian Cox processes that allows for the inclusion of study-specific variables. We present an efficient MCMC scheme based on the Hamiltonian Monte Carlo algorithm to simulate draws from the posterior. Computational time is significantly reduced through a parallel implementation using a graphical processing unit card. We evaluate the method on both simulated and real data.

1 Introduction

Functional magnetic resonance imaging (fMRI) has experienced an explosive growth over the past two decades, leading to significant advances in our understanding of human brain function (Raichle, 2003). The objective of an fMRI study is to identify the neural correlates of a physical, mental, or perceptual process. During the experiment, participants are asked to perform a series of tasks, such as viewing images or reading text, while the MRI device

measures changes in the Blood Oxygenation Level Dependent (BOLD) signal, a surrogate for neuronal activity. A series of 3D volumes of BOLD measurements are collected for each participant and at each *voxel* (volume element) a hierarchical multi-subject time series model is fit. The ultimate result of this process is a statistic image that assesses the evidence for a change in the BOLD signal induced by the task. For more details on the statistical analysis of fMRI data see Lindquist (2008) or Poldrack *et al.* (2011).

While fMRI is now a standard tool, the importance of the findings can be undermined by the inherent limitations of a single fMRI experiment. In particular, there are three facets of fMRI experiments that put the quality of the results into question. First, fMRI studies suffer from low power. The typical sample size of an fMRI study is small, with the majority of experiments using less than 20 participants (Carp, 2012). For example, a recent meta-analysis of emotions found a median sample size of 13 (Lindquist *et al.*, 2012). Such small sample sizes limit power and surely result in Type II errors. Furthermore, Type I errors are also prevalent (Wager *et al.*, 2009), as a result of searching for significant effects over thousands of voxels. While methods exist to correct multiplicity, sometimes they either are not applied or are applied inappropriately. Finally, fMRI experiments suffer from low reliability (Raemaekers *et al.*, 2007); that is, an individual examined under the same experimental conditions may give different results on two different sessions.

For all these reasons, there has been a growing interest in meta-analysis of functional neuroimaging studies. A meta-analysis can increase power by combining the available information across studies, hence separating the consistent findings from those occurring by chance. In fMRI there are two broad approaches for meta-analysis. When the full statistical images are available, that is effect sizes and associated standard errors for all voxels in the brain, an *Intensity-Based Meta-Analysis* (IBMA) can proceed by means of standard meta-analytic mathods (see Hartung *et al.* (2008) for an overview). However, these statistic images (200,000+ voxels) traditionally have not been shared by authors. Instead, researchers only publish the x, y, z brain atlas coordinates of the local maxima in significant regions of the statistic image. We call these coordinates the *foci* (singular *focus*). When only foci are available then a *Coordinate-Based Meta-Analysis* (CBMA) is conducted. As can be expected, the transition from full images to the lists of foci involves a heavy loss of information (Salimi-Khorshidi *et al.*, 2009). However, since the vast majority of researchers rarely provide the full images, CBMA constitutes the main approach for fMRI meta-analysis.

Most work in the field is focused on the so-called kernel-based methods such as *Activation Likelihood Estimation* (Turkeltaub *et al.*, 2002; Eickhoff *et al.*, 2012, ALE), *Multilevel Kernel Density Analysis* (Wager *et al.*, 2004, 2007, MKDA) and *Signed Differential Mapping* (Radua and Mataix-Cols, 2009; Radua *et al.*, 2012, SDM). Briefly, these methods construct a statistic map as the convolution of the foci[1] with 3D spatial kernels. Areas

[1]Precisely, this is a convolution of a Dirac delta function located at each foci with a given kernel.

of the map with large values suggest brain regions of consistent activation across studies. For statistical inference, the map is thresholded by reference to a Monte Carlo distribution under the null hypothesis of uniform spatial foci allocation. Kernel-based methods are not based on an explicit probabilistic model and hence often lack interpretability. Moreover, for some methods it is difficult to obtain standard errors and hence only p-values are reported for each voxel. Recently, some model-based methods were proposed to address the limitations of kernel-based methods, such as the *Bayesian Hierarchical Independent Cox Cluster Process* model of Kang *et al.* (2011) and the *Bayesian Nonparametric Binary Regression* model of Yue *et al.* (2012). However, none of these methods account for study-specific characteristics in the meta-analysis.

In this work we propose a Bayesian spatial point process model, an extension of the log-Gaussian Cox process model (Møller *et al.*, 1998) that can include study specific characteristics as explanatory variables, thus introducing the notion of *meta-regression* (Greenland, 1994) to CBMA. Meta-regression is an important facet of meta-analysis, especially when there is appreciable heterogeneity among the studies. For example, in a review of 241 fMRI studies, Carp observed 223 different analytical strategies. It is therefore essential to explore the effect that study characteristics have on the analyses outcomes, as well as to identify which of the outcomes replicate regardless of the analytical pipelines. Further, our method takes advantage of some recent advances in parallel computing hardware, specifically the use of graphical processing units (GPUs).

The remainder of this manuscript is organised as follows. In Section 2 we outline the details of our model and describe the approximations required for posterior inferences. In Section 3 we describe the Hamiltonian Monte Carlo algorithm used for simulating samples from the posterior distribution of our model parameters. In Sections 4 and 5 we apply the model to simulated and real data respectively. In Section 6 we conclude and set some possible directions for future work.

2 Model specifications

Suppose that there are a total number of I studies in the meta analysis and that each study i comes with a point pattern $\mathbf{x}_i$, a set of foci $x_{ij} \in \mathcal{B} \subset \mathbb{R}^3$, where $\mathcal{B}$ is the support of the analysis, usually set from a standard atlas of the brain, and $j = 1, \ldots, n_i$, where n_i is the number of foci in a study. Additionally, suppose that for each point pattern there is a set of K study specific characteristics, $\{z_{ik}\}_{k=1}^{K}$.

We will assume that each point pattern $\mathbf{x}_i$ is the realisation of a Cox point process $\mathbf{X}_i$ defined on $\mathcal{B}$, driven by a random intensity $\lambda_i(\cdot)$. We can then model the intensity function at each point $\xi \in \mathcal{B}$ as:

$$\log \lambda_i(\xi) = \beta_0(\xi) + \sum_{k=1}^{K} \beta_k(\xi) z_{ik}, \tag{1}$$

where β_0 is the baseline intensity parameter, and the β_k are the regression coefficients. Equation (1) defines a spatial log-linear model over the brain. Foci are more likely to occur in regions of the brain with high intensity values whereas we expect almost no foci in regions as the intensity approaches zero. The exact rates are given by the properties of a Cox process. In particular, given $\lambda_i(\cdot)$, the expected number of foci in any bounded $B \subseteq \mathcal{B}$ is a Poisson random variable with mean $\int_B \lambda_i(\xi)d\xi$ (Møller and Waagepetersen, 2004).

In practice we expect that some covariates will have a global (homogenous) effect. Therefore, we split the covariates into the $K^* \leq K$ that have a local effect and $K - K^*$ that have a global effect:

$$\log \lambda_i(\xi) = \sum_{k=0}^{K^*} \beta_k(\xi) z_{ik} + \sum_{k=K^*+1}^{K} \beta_k z_{ik}, \tag{2}$$

where $z_{i0} = 1$ for all $i = 1, \ldots, I$.

A Bayesian model is defined with prior distributions on model parameters, here the functional parameters $\beta_k(\cdot)$, $k = 0, \ldots, K^*$, and scalar parameters β_k, $k = K^* + 1, \ldots, K$. A natural way to proceed is to assume that $\beta_k(\cdot)$ are realisations of Gaussian processes and that the β_k have normal distributions. That way, the right hand side of Equation (2) is also a Gaussian process, and each point process is a log-Gaussian Cox process (LGCP) (Møller *et al.*, 1998). We will assume an isotropic, power exponential correlation structure, that is for points $\xi, \xi' \in \mathcal{B}$ we have:

$$\mathbb{C}\text{or}\left(\boldsymbol{\beta}(\xi), \boldsymbol{\beta}(\xi')\right) = \exp\left\{-\rho_k ||\xi - \xi'||^{\delta_k}\right\}, \tag{3}$$

for $k = 1, \ldots, K^*$, where $\rho_k > 0$ are the correlation decay parameters and $\delta_k > 0$ are the smoothness parameters. The same correlation structure was used by Møller *et al.* (1998) and Møller and Waagepetersen (2003) in the context of LGCPs. Of course, one could consider alternative correlation structures such as the *Matérn* covariance function (see for example Rasmussen and Williams (2005)).

The log-Gaussian Cox process is a flexible model for spatial point data that can account for aggregation (Møller *et al.*, 1998; Møller and Waagepetersen, 2007) or even repulsion between points (Illian *et al.*, 2012a) and has therefore found applications in several fields such as disease mapping (Benes *et al.*, 2002; Liang *et al.*, 2009) and ecology (Møller and Waagepetersen, 2003; Illian *et al.*, 2012b).

By the definition of a Cox process, $\mathbf{X}_i$ is a Poisson point process on $\mathcal{B}$ conditional on $\lambda_i(\cdot)$ (Møller and Waagepetersen, 2004). The density (Radon-Nikodym derivative) of this

point process with respect to the unit rate Poisson process is:

$$\pi\left(\mathbf{x}_i \mid \lambda_i\right) = \exp\left\{|\mathcal{B}| - \int_{\mathcal{B}} \lambda_i(\xi) d\xi\right\} \prod_{x_{ij} \in \mathbf{x}_i} \lambda_i(x_{ij}), \tag{4}$$

for $i = 1, \ldots, I$, with $|\mathcal{B}|$ denoting the volume of the brain. We can view $\pi\left(\mathbf{x}_i \mid \lambda_i\right)$ as the density of the sampling distribution of the data; if we further assume independent studies, then we obtain the likelihood for the model as:

$$\mathcal{L}\left(\{\beta_k(\cdot)\}_{k=0}^{K^*}, \{\beta_k\}_{k=K^*+1}^{K}\right) = \prod_{i=1}^{I} \pi\left(\mathbf{x}_i \mid \lambda_i\right), \tag{5}$$

where λ_i is as defined in Equation (2). Inference can be then achieved through the posterior distribution of the model which is given, up to a normalising constant, by:

$$\pi\left(\{\beta_k(\cdot)\}_{k=0}^{K^*}, \{\beta_k\}_{k=K^*+1}^{K} \mid \{\mathbf{x}_i\}_{i=1}^{I}\right) \propto \mathcal{L}\left(\{\beta_k(\cdot)\}_{k=0}^{K^*}, \{\beta_k\}_{k=K^*+1}^{K}\right) \prod_{i=1}^{K^*} \pi(\beta_k(\cdot)) \prod_{k=K^*+1}^{K} \pi(\beta_k), \tag{6}$$

where $\pi(\beta_k(\cdot))$ and $\pi(\beta_k)$ are the priors on the functional and scalar parameters, respectively. We discuss the priors below in Section 2.1.

2.1 Posterior approximation

Calculation of the posterior in Equation (6) requires the evaluation of the infinite dimensional Gaussian processes $\beta_k(\cdot)$, $k = 0, \ldots, K^*$, which we approximate with a finite dimensional distribution. Following Møller *et al.* (1998) and Benes *et al.* (2002), we consider the discretisation of the 3D volume with a regular rectangular grid $W \supset \mathcal{B}$. We use V cubic cells (i.e. voxels) in W with volume $A = a^3$, where a is the length of the side. In neuroimaging, analysis with 2mm^3 cubic voxels is typical, leading to a box-shaper grid of about 1 million voxels, of which about 200,000 are in the brain or cerebellum. Note that for simplicity, we consider both grey matter and white matter voxels in our implementations. Voxels are indexed $v = 1, \ldots, V$, and the coordinate of v is the location of the center $\boldsymbol{\nu}_v \in \mathbb{R}^3$.

For any $k = 0, \ldots, K^*$, the Gaussian process $\beta_k(\cdot)$ can be now approximated with a step function which is constant within each voxel v and equal to the value of $\beta_k(\cdot)$ at the location of the center, i.e. $\beta_k(\boldsymbol{\nu}_v)$. Waagepetersen (2004) shows that the accuracy of this approximation improves as a goes to zero. By definition, $\boldsymbol{\beta}_k = [\beta_k(\boldsymbol{\nu}_1), \ldots, \beta_k(\boldsymbol{\nu}_V)]$ are multivariate Gaussian vectors. We parametrise $\boldsymbol{\beta}_k$ as:

$$\boldsymbol{\beta}_k = \mu_k + \sigma_k \mathbf{R}_k^{1/2} \boldsymbol{\gamma}_k, \tag{7}$$

where μ_k are the overall (scalar) means, σ_k are the marginal standard deviations, $\mathbf{R}_k$ are the $V \times V$ correlation matrices with elements $(\mathbf{R}_k)_{ij} = \exp\left\{-\rho_k ||\boldsymbol{\nu}_i, \boldsymbol{\nu}_j||^{\delta_k}\right\}$, and $\boldsymbol{\gamma}_k$ are the *a priori* $\mathcal{N}_V(\mathbf{0}, \mathbf{I}_V)$ vectors, $k = 0, \ldots, K^*$. The same parametrisation is used by Møller *et al.* (1998), Christensen and Waagepetersen (2002) and is advocated by Christensen *et al.* (2006) because it allows for computationally efficient posterior simulations. For the purposes of this work we will only consider the case where $\delta_k = 2$ i.e. the *Gaussian correlation* function. This choice may seem rather simplistic but is justified by the sparsity of points in CBMA data and ubiquitous use of Gaussian smoothing kernels in neuroimaging data analysis.

Priors for the V-vectors $\boldsymbol{\gamma}_k$ are induced by the parametrisation of Equation (7). For the remaining parameters we will assume weekly informative priors: $\mu_k \sim \mathcal{N}(0, 10^8)$, $\sigma_k \sim \mathcal{N}(0, 10^8)$, $\beta_k \sim \mathcal{N}(0, 10^8)$ and $\rho_k \sim \text{Uni}[0, 100]$[2].

Once the latent Gaussian processes are approximated, one can also approximate λ_i with a step function as before. The intensities at the center of each voxel are given by:

$$\boldsymbol{\lambda}_i = \exp\left\{\sum_{k=0}^{K^*}\left(\mu_k + \sigma_k \mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right) z_{ik} + \sum_{i=K^*+1}^{K} \beta_k z_{ik}\right\}, \tag{8}$$

where $\boldsymbol{\lambda}_i$ is the V-vector, the discretised intensity. We will write $\lambda_{iv} = (\boldsymbol{\lambda}_i)_v$ for the v-element of study i's intensity, and note we require $z_{i0} = 1$ to capture the mean effect. The approximated posterior is:

$$\pi\left(\boldsymbol{\theta} \mid \{\mathbf{x}_i\}_{i=1}^{I}\right) \propto \prod_{i=1}^{I}\left[\exp\left\{-\sum_v A_v \lambda_{iv}\right\}\prod_{j=1}^{n_i}\lambda_{iv(x_{ij})}\right]\pi(\boldsymbol{\theta}), \tag{9}$$

where $\boldsymbol{\theta} = \left\{\{\mu_k\}_{k=1}^{K^*}, \{\sigma_k\}_{k=1}^{K^*}, \{\rho_k\}_{k=1}^{K^*}, \{\boldsymbol{\gamma}_k\}_{k=1}^{K^*}, \{\beta_k\}_{k=K^*+1}^{K}\right\}$, A_v takes on the value A when $\boldsymbol{\nu}_v \in \mathcal{B}$ and 0 otherwise, $v(x_{ij})$ is the index of the voxel containing x_{ij}, and $\pi(\boldsymbol{\theta})$ is the joint prior distribution of the parameters. The posterior distribution in Equation (9) is still analytically intractable due to the presence of an unknown normalising constant and thus we need to resort to Monte Carlo simulation or approximation techniques to obtain samples from it.

3 Sampling algorithm details

Bayesian methodology for inference on LGCPs can be broadly divided into two main categories: simulation based approximations of the posterior such as Markov Chain Monte Carlo (Møller *et al.*, 1998) and *elliptical slice sampling* (Murray *et al.*, 2010), and deterministic approximations to the posterior such as *integrated nested Laplace approximations* (Illian *et al.*, 2012a; Simpson *et al.*, 2016, INLA) and *variational Bayes* (Jaakkola and Jor-

[2]This prior may not be uninformative for applications with a different scale of distances

dan, 2000). In a recent study, Taylor and Diggle (2014) compare the Metropolis-adjusted Langevin (MALA) algorithm with INLA and find that both methods give similar results. In our application, we choose to use simulation based methods because application on our 3D problem is more straightforward.

Of course, building an algorithm for such high-dimensional problem is non-trivial. Girolami and Calderhead (2011) showed that of all possible strategies, their Riemann Manifold Hamiltonian Monte Carlo (RMHMC) sampler is the computationally most efficient for LGCPs in a 2D setting. Unfortunately, application in this problem (3D setting) is prohibitive as it would require the inversion of a huge $V \times V$ matrix. Therefore, we choose to use the standard Hamiltonian Monte Carlo (Duane *et al.*, 1987; Neal, 2011, HMC) algorithm which Girolami and Calderhead (2011) found to be the second most efficient.

HMC initially appeared in the physics literature by Duane *et al.* (1987) under the name *Hybrid Monte Carlo*, and later emerged into statistics literature by Neal (2011). HMC emulates the evolution of a particle system which is characterised by its position ($\mathbf{q}$) and momentum ($\mathbf{p}$) over time. In our case, $\mathbf{q}$ will be the parameter vector of interest $\boldsymbol{\theta}$ and $\mathbf{p}$ will be introduced artificially from a $\mathcal{N}_d(0, \mathbf{M})$ distribution, with d being the dimensionality of the problem and $\mathbf{M}$ the mass matrix. The dynamics of the system are described by a set of differential equations, known as Hamilton's equations.

HMC alternates between moves for the position vector $\boldsymbol{\theta}$ and the momentum vector $\mathbf{p}$ based on Hamilton's equations. If the solutions of the equations can be found analytically then moves will be deterministic; if not, numerical integration is required and an acceptance/rejection step must be performed to account for integration error. Integration is done in fictitious time ϵL, where ϵ is the *stepsize* and L is the *number of steps*. Typically the *leapfrog integrator* is employed, which for $L = 1$ and starting at time t is performed as (Neal, 2011):

$$\begin{aligned} \mathbf{p}\left(t + \frac{\epsilon}{2}\right) &= \mathbf{p}\left(t\right) + \frac{\epsilon}{2}\nabla_{\boldsymbol{\theta}} \log \pi\left(\boldsymbol{\theta}(t) \mid \{\mathbf{x}_i\}_{i=1}^{I}\right) \\ \boldsymbol{\theta}\left(t + \epsilon\right) &= \boldsymbol{\theta}\left(t\right) + \epsilon \mathbf{M}^{-1}\mathbf{p}\left(t + \frac{\epsilon}{2}\right) \\ \mathbf{p}\left(t + \epsilon\right) &= \mathbf{p}\left(t + \frac{\epsilon}{2}\right) + \frac{\epsilon}{2}\nabla_{\boldsymbol{\theta}} \log \pi\left(\boldsymbol{\theta}(t + \epsilon) \mid \{\mathbf{x}_i\}_{i=1}^{I}\right). \end{aligned} \tag{10}$$

Overall, if the method is applied correctly, it will produce a time-reversible Markov chain that has the desired distribution $\pi\left(\boldsymbol{\theta} \mid \{\mathbf{x}_i\}_{i=1}^{I}\right)$ as its stationary distribution. As we show in Appendix A, gradient expressions are available in closed form for all model parameters including correlation parameters ρ_k. We therefore choose to update all $\{\mu_k, \sigma_k, \rho_k, \boldsymbol{\gamma}_k\}_{k=0}^{K^*}$ and $\{\beta_k\}_{k=K^*+1}^{K}$ jointly with HMC. The solutions to Hamilton's equations are not available analytically so we need to use the Leapfrog integrator and include an accept/reject step at the end of it.

The procedure requires the specification of a stepsize ϵ and a total number of leapfrog steps L. Hoffman and Gelman (2014) show how tuning can be achieved automatically but

when we applied this method to our problem running time was increased substantially. Therefore we use an alternative approach to tune these parameters. The stepsize is automatically adjusted during the burn-in phase of the HMC to give an overall acceptance rate close to the 65% suggested by Neal (2011). In particular, if ϵ_t is the stepsize at iteration t and q_{t_1} is the acceptance rate over the past t_1 iterations, then every t_2 iterations we calculate the new stepsize ϵ'_t as:

$$\epsilon'_t = \begin{cases} 0.9\epsilon_t & q_{t_1} < 0.60 \\ \epsilon_t & 0.60 \leq q_{t_1} \leq 0.70 \\ 1.1\epsilon_t & q_{t_1} > 0.70 \end{cases}. \tag{11}$$

Specifically we use $t_1 = 100$ and $t_2 = 10$ A similar approach is empoyed by Marshall and Roberts (2012) for MALA. The latter (number of leapfrog steps), is always fixed to $L = 50$. We took this approach because we found that, for our LGCP application, the mixing properties of the algorithm scale linearly with L but also with the total number of HMC iterations. Hence one can use a relatively large L and few iterations or relatively smaller L and more iterations, the total computation time staying relatively constant.

The last tuning parameter in the HMC algorithm is the variance-covariance matrix of the zero mean normal momentum parameters, $\mathbf{M}$. To our knowledge, there is only limited off the shelf methodology on how to adjust $\mathbf{M}$. As a starting place we set $\mathbf{M} = \mathbf{I}$. Neal (1996) suggests that if an estimate of the posterior variance $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}$ is available then a good practice is to set $\mathbf{M} = \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}^{-1}$. In principle, $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}$ can be estimated during the burn-in phase of HMC but in practice this is not possible due to the dimensionality of the problem. In our simulations, we found that the mean posterior variance of the elements of the $\boldsymbol{\gamma}_k$ was higher compared to the scalar parameters, followed by β_k or σ_k and then ρ_k. Especially for the ρ_k the scale is typically much smaller compared to the other parameters in our applications and so we use $100 \times \rho_k$ instead of ρ_k. After the reparametrisation we found that setting the mass for parameters of $\boldsymbol{\gamma}_k$, β_k, σ_k and ρ_k equal to 1, 3, 3 and 10 respectively worked well in all implementations.

The most computationally demanding part of the algorithm is the the calculation of the large matrix-vector products $\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k$ appearing in the intensity functions of Equation (8). Luckily, an elegant solution to this problem is given by Møller *et al.* (1998) based on *circulant embedding* that was first proposed by Dietrich and Newsam (1993) and Wood and Chan (1994). The key to the approach is the linear algebra result that a circulant matrix has the discrete Fourier basis as its eigenvectors. $\mathbf{R_k}$ is not circulant but is block toeplitz and can be embeded in a $(2V) \times (2V)$ matrix that is circulant. Thus the matrix square root, inversion and multiplication can be accelerated by using (the highly efficient) discrete Fourier transform (DFT) of the embedded matrix and manipulating Fourier coefficients, followed by inverse DFT and extracting the appropriate sub-matrix/sub-vector. See Rue

and Held (2005) for more details.

We close this section by stressing that despite the massive dimensionality of the parameter vector, the problem has a very high degree of parallelisation. Intensities can be evaluated in blocks of thousands of voxels simultaneously making the algorithm suitable for implementation in a graphical processing unit (GPU). The most computationally intensive part of our model, namely operations with DFTs, is also amenable to parallelisation and there exist libraries such as NVIDIA's cuFFT library that are designed for this specific task. Overall, we believe that implementation of the log-Gaussian Cox process model described above will soon become a routine task for any moderately powerful GPU device.

4 Simulation studies

We consider two simulation setups. In the first we draw samples directly from the log-Gaussian Cox process whereas in the second we create synthetic studies based on a different model to assess the robustness of our method to model misspecification. For consistency, all processes are defined on the same brain atlas used in the application of Section 5, consisting of $216,040$ 2mm^3 cubic voxels. The average number of foci per simulated study is kept low (mean number of foci per study is 5) to resemble the sparsity of points observed in real CBMA data. Finally, the total number of studies is fixed to 200 in both analyses, similar to the sample sizes available in real applications (Kang *et al.*, 2011, for example).

4.1 Setup 1

In this setting we simulate 200 studies, with two spatially varying covariates that account for the mean of two groups of studies, and two non-spatially varying covariates. For $i = 1, \ldots, 200$ we set:

$$\lambda_{iv} = \exp\left\{\sum_{k=1}^{2}\left(\mu_k + \sigma_k\left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_v\right)z_{ik} + \sum_{i=3}^{4}\beta_k z_{ik}\right\}, \tag{12}$$

where $z_{i1} \sim \text{Bernoulli}(0.5)$, $z_{i2} = 1 - z_{i1}$, $z_{i3} \sim \text{Uniform}[-1, 1]$ and $z_{i4} \sim \text{Bernoulli}(0.5)$. Note that this parametrisation of the covariates implies existence of two types of studies, say type 1 and 2, with different spatially varying means and the effect of one continuous and one categorical covariate. The expected total number of foci is 3.99 and 4.16 for studies of type 1 and 2 respectively. We draw $\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2$ from their $\mathcal{N}_V(0, \mathbf{I})$ prior and fix the values of the scalar parameters shown in Table 1. We run the HMC algorithm of Section 3 for 10,000 iterations, discarding the first 4,000 as a burn-in and save every 6 iterations for a total of 1,000 saved posterior samples. This took roughly 14 hours on an NVIDIA Tesla K20c GPU card.

Results are summarised in Table 1 and Figure 1. In Table 1 we see that the scalar

parameters are estimated accurately despite the sparsity of points in the realisations. The 95% credible intervals contain the true values of all the parameters in the setup. Some traceplots for these parameters can be found in Section B of the Appendix.

For $z_{i3} = z_{i4} = 0$, the median expected number of points is 3.97 (95% CI [3.84,4.10])for type 1 and 4.61 (95% CI [4.46,4.78]) for type 2. These values are very similar to the values we observe in the simulated dataset, that is 3.98 for type 1 and 4.53 for type 2. This indicates that our model does a good job fitting the data. The shape of the latent Gaussian processes $\mu_k + \sigma_k \mathbf{R}_k^{1/2} \boldsymbol{\gamma}_k$ is generally captured for both types as can be seen in Figure 1. In particular, we can see that the maxima in the true and estimated images appear roughly in the same locations. The same cannot be said about the other values but this is expected given the dearth of information in regions of low intensity.

Table 1: Posterior summaries of the scalar parameters of the LGCP model, fit to the simulated data of Section 4.1. Results are based on 1,000 posterior draws. The values for the correlation parameters ρ_1, ρ_2 are multiplied by 100. The values for β_1 and β_2 are multiplied by 10.

Parameter	**True Value**	**Posterior median**	95% **credible interval**
μ_1	-13.7	-13.72	[-13.99 , -13.48]
μ_2	-14.2	-14.14	[-14.47 , -13.86]
σ_1	1.2	1.19	[1.01 , 1.38]
σ_2	1.6	1.61	[1.43 , 1.81]
ρ_1	1	0.93	[0.69 , 1.27]
ρ_2	2	2.30	[1.69 , 3.15]
β_3	2	1.44	[0.22 , 2.52]
β_4	1	0.95	[0.32 , 1.65]

4.2 Setup 2

In this setup we create datasets with a pattern of points that follows brain structures of interest. Again there are two types of studies, say type 1 and type 2. For each study i, $i = 1, \ldots, 200$, we generate the total number of points from a Negative Binomial distribution with mean $\mu = 6 + 2z_{i3} - \mathbf{1}_{\{z_{i4}=0\}} + \mathbf{1}_{\{z_{i4}=1\}}$ and variance $\mu^2/20$. For the covariates, $z_{i3} \sim \text{Uni}[-1, 1]$ and $z_{i4} \sim \text{Bernoulli}(0.5)$. Once we know the exact number of foci per study, we assign the study uniformly at random to one of the 2 types and the distribute its foci as follows. For type 1, foci appear systematically in the following regions: each focus can be observed in the right amygdala (B_R) with probability 55%, the orbifrontal cortex (B_C) with probability 30% or anywhere else in the brain with probability 15%. The configuration for type 2 differs in that most of the points will go to the left amygdala (B_L) instead of the right amygdala. If a focus is assigned to one of the three broad regions, the exact location has a uniform distribution over the region. In the fourth column of Figure

2 the regions in red and blue correspond to the left and right amygdala respectively while the orbifrontal cortex is coloured in green.

HMC is run for 10,000 iterations, discarding the 4,000 first as a burn-in and saving every 6 to obtain a total of 1,000 samples from the posterior. The run took approximately 15 hours on a Tesla K20c GPU card.

Results are shown in Figure 2 where in the first two columns we see median posterior log-intensities for the two types, in different axial slices. In both cases, we find that the regions with the highest intensities are the amygdalae and that the orbifrontal cortex is a region of high posterior intensity as well. The median expected number of points is 5.81 for type 1 (95% CI [5.36,6,32]) and 6.45 for type 2 (95% CI [5.97,6.97]). The observed values are 6.27 and 6.73 respectively.

Conditional on there being exactly one focus, we can estimate the probability that this focus appears in any subset $B \subseteq \mathcal{B}$ as $\int_B \lambda(\xi)\, d\xi / \int_{\mathcal{B}} \lambda(\xi)\, d\xi$. Using the posterior draws obtained from the HMC algorithm, we can obtain the posterior distribution of any such quantity. For our simulated type 1 data we find that the median posterior probability of observing a focus in the right amygdala (B_R) is 0.43 (95% CI [0.40,0.48]). For type 2, the probability of observing a focus in the left amygdala (B_L) is 0.42 (95% CI [0.39,0.46]). For the orbifrontal cortex (B_C) the median posterior probabilities are 0.25 for type 1 and 0.23 for type 2, with 95% credible intervals $[0.22, 0.28]$ and $[0.20, 0.26]$ respectively. We therefore see that the model underestimates the probabilities for B_R, B_L and B_C. This bias can be attributed to the smoothness that is imposed by our parameter δ thus leading to increases intensities just outside these regions as well as regions where noise foci appear.

An interesting question one may ask is which are the regions of the brain that are activated by one type or the other, but not both. To answer this, one can construct the mean standardised posterior difference map computed as the ratio of the posterior mean of the difference $(\boldsymbol{\beta}_1)_v - (\boldsymbol{\beta}_2)_v$, to the posterior standard deviation of that difference: $\frac{(\boldsymbol{\beta}_1)_v - (\boldsymbol{\beta}_2)_v}{\mathrm{sd}\left((\boldsymbol{\beta}_1)_v - (\boldsymbol{\beta}_2)_v\right)}$. Extreme negative or positive values are evidence of differences between the two types. We show the difference map in the third column of Figure 2. As we see, the model distinguishes the the two types in the amygdala but the differences are small in the rest of the brain. This is a very interesting feature of the model, especially for applications in CBMA where researchers are sometimes interested in comparing a similar process in a different domain, see for example Rottschy *et al.* (2012) or Section 5 for an application. An alternative way to do the comparison would be to use the posterior intensity draws to find $\mathbb{P}\left(\boldsymbol{\lambda}_v^1 - \boldsymbol{\lambda}_v^2 > \epsilon\right) \approx \frac{\sum_{n=1}^{N} \mathbf{1}_{\{\lambda_v^{1n} - \lambda_v^{2n} > \epsilon\}}}{N}$, where ϵ is a threshold difference, λ_v^{kn} is the posterior intensity for type k in voxel v obtained from the n-th iteration of the algorithm, and N is the total number of HMC draws. However, in this approach it is hard to choose the threshold difference ϵ.

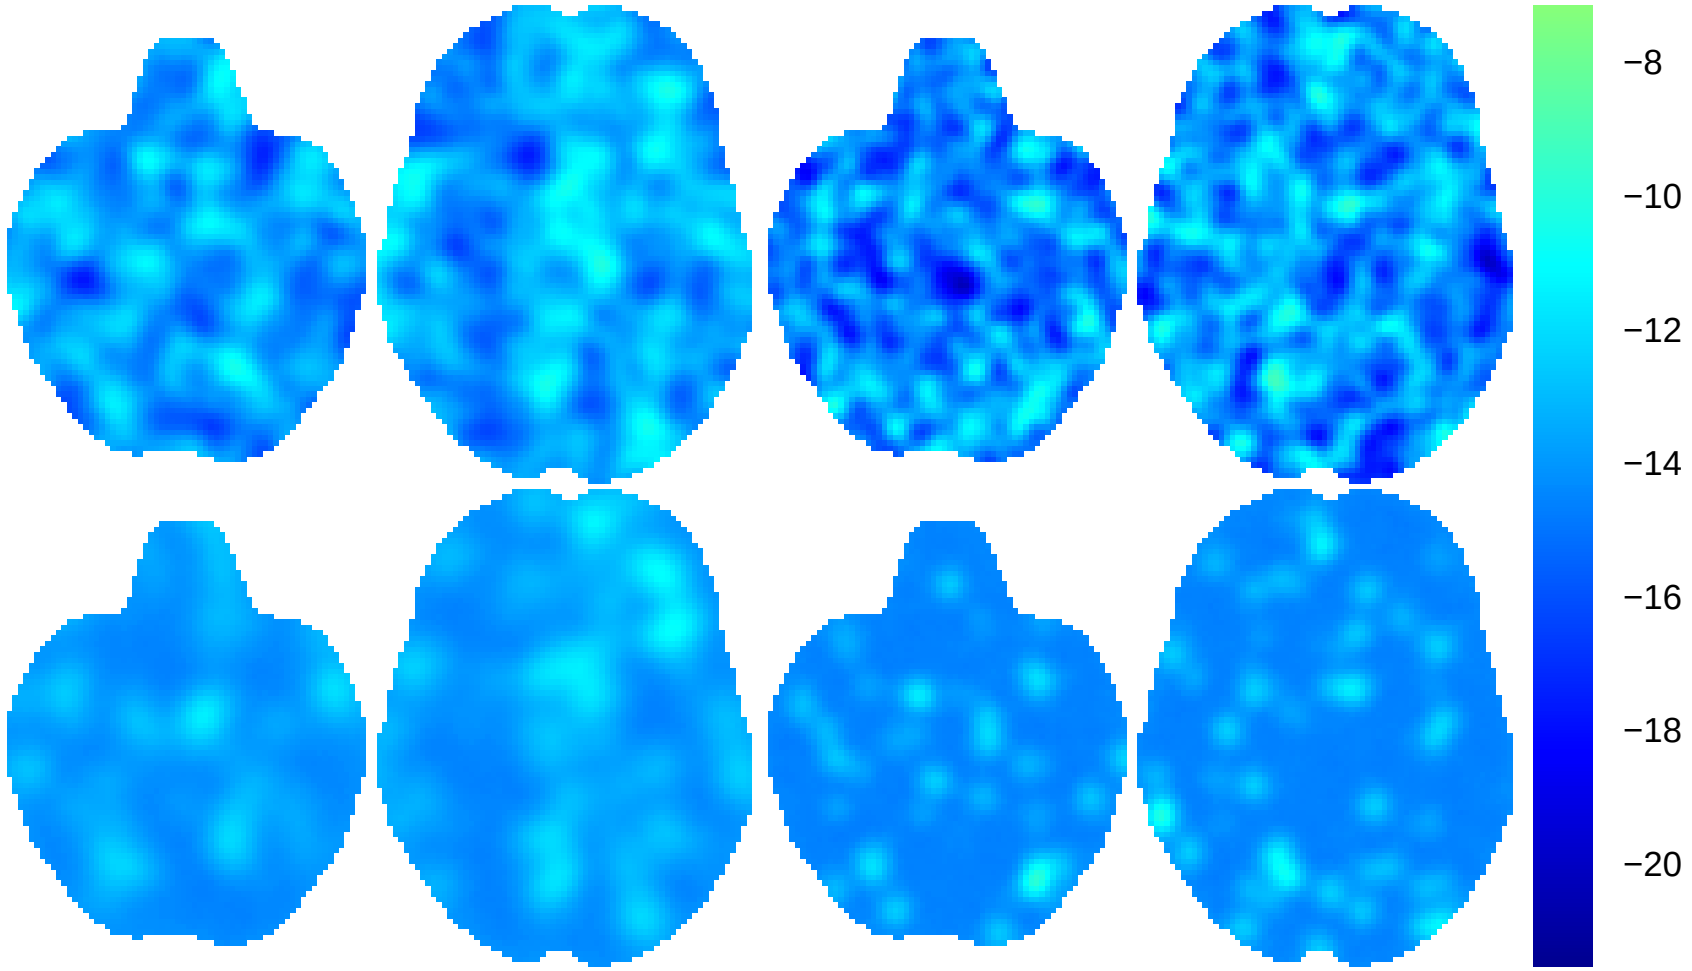

Figure 1: Some true (top row) and estimated (bottom row) latent Gaussian processes for type 1 (columns 1 and 2) and type 2 (columns 3 and 4) in the simulation setup 1 of Section 4.1. Columns 1 and 3 correspond to axial slice $z = -22$; columns 2 and 4 correspond to axial slice $z = 4$. While they may appear dissimilar at first, observe that the most intense regions of the true and estimated intensity match up; in this punishingly sparse setting (mean ≈ 4 foci per 3D image), the less intense regions have too few points to learn the intensity.

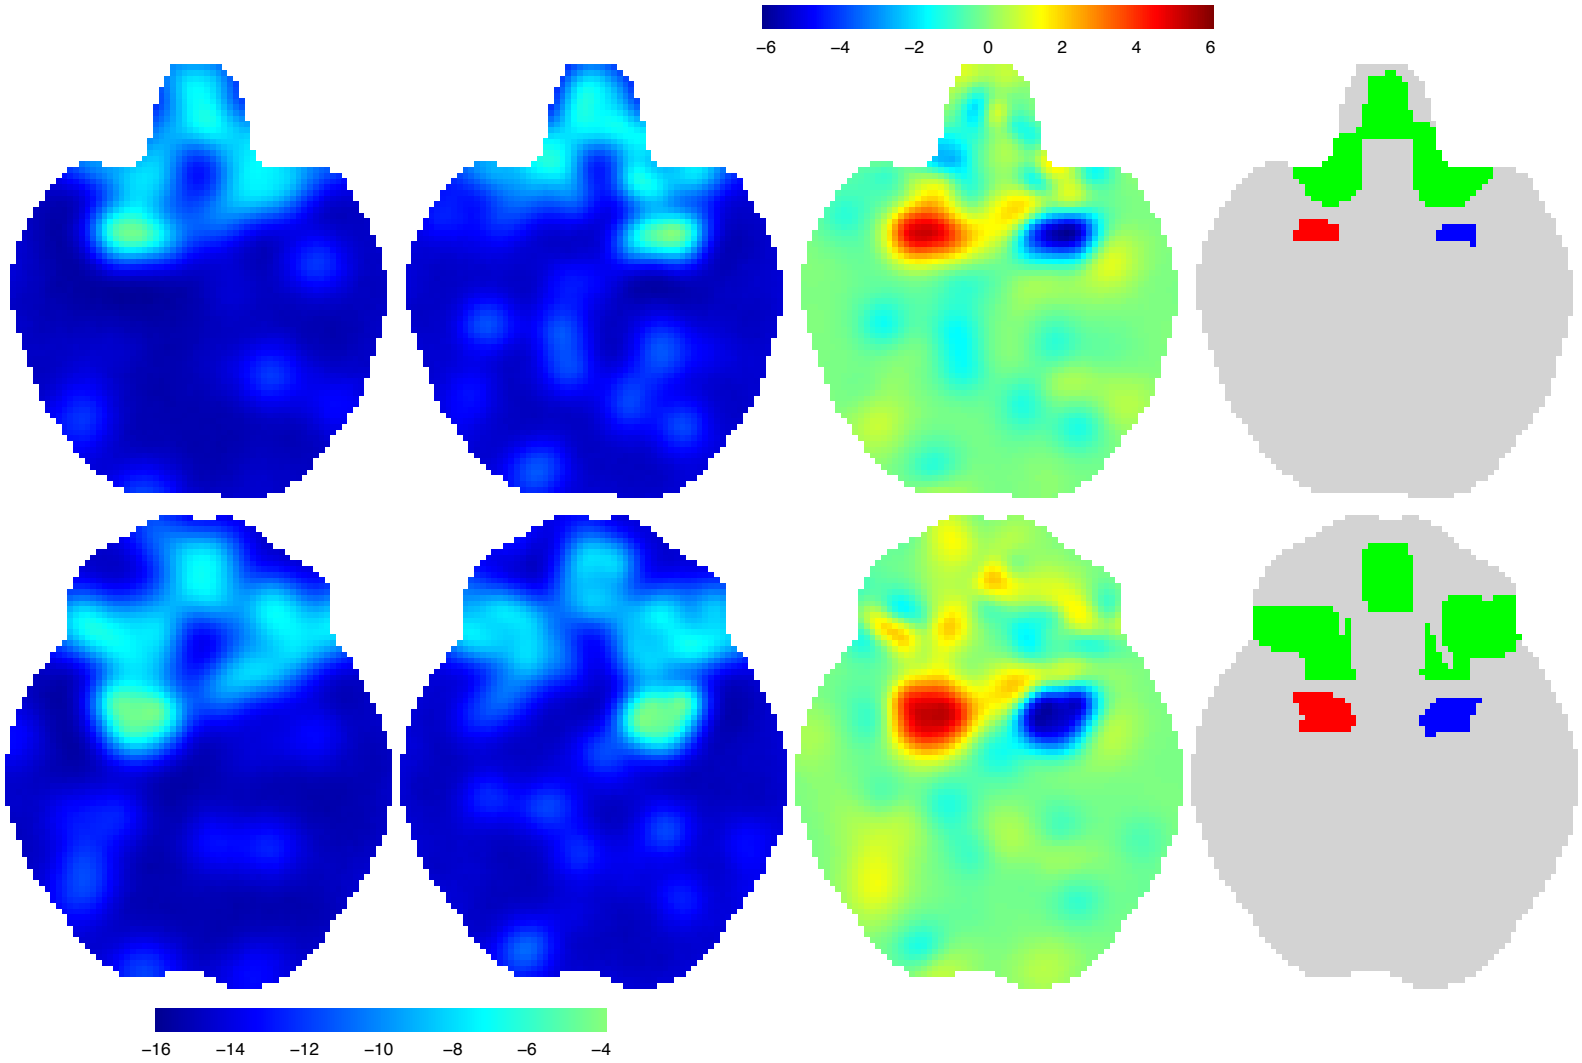

Figure 2: Results for simulation setup 2 of Section 4.2. The top row corresponds to axial slice $z = -24$ whereas the bottom row corresponds to axial slice $z = -16$. Columns 1 and 2 are the estimated log-intensities for type 1 and type 2 respectively. The third column is the standardised mean posterior difference between the two latent Gaussian processes in the corresponding slice; bright colours indicate areas mostly activated by type 1 process. The fourth column shows the regions of the brain systematically activated by the two processes; red for type 1, blue for type 2 and green for both.

5 Application: meta-analysis of emotion and executive control studies

In this Section we apply our model to a real meta-analysis data set.

5.1 Data despcription

Our dataset consists of 1,193 neuroimaging studies. The studies were conducted between 1985 and 2015 and the average number of participants was 16. Of these, 855 studies are on emotion and the remaining 338 are on executive control. The sample has a total of 10,266 foci, 6,112 from emotion (7.15 on average) and 4,154 from executive control (12.29 on average). Figure 3 is a graphical representation of the data. We see that even though there is some clustering, the foci are distributed throughout the brain. Our application will focus on identifying the regions of consistent activation across studies and infer on possible difference between the two types.

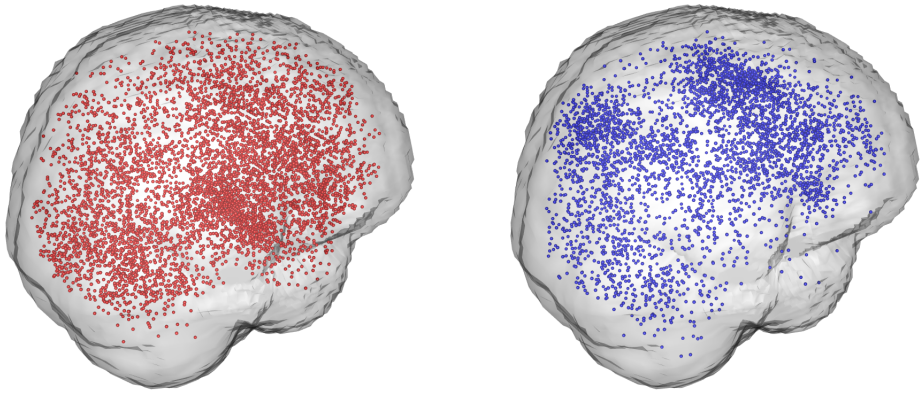

Figure 3: Graphical representation of the meta-analysis dataset. Data consist of foci from 1,193 studies of two types, 855 emotion (left, red) and 338 exectutive control (right, blue) studies. The overall number of foci is 6,112 and 4,154 for emotion and executive control respectively, for a combined total of 10,266 foci. The code used to generate the figure is courtesy of Jian Kang.

5.2 Algorithm details and convergence diagnostics

We use the same parametrisation as in Section 4: we have 2 spatially varying intercepts, one for emotion and one for executive control. We run the HMC for 15,000 iterations in total, discarding the first 5,000 as a burn-in. The total number of leapfrog steps is set to $L = 50$ and the stepsize is initialised at $\epsilon = 0.0001$. We use a diagonal mass matrix with units specified as explained in Section 3.

Convergence of the HMC chain is assessed visually by inspection of posterior traceplots for the model parameters. We run a total of 3 HMC chains in order to examine if they

all converge to the same values. Posterior traceplots, along with autocorrelation plots are shown in Section C of the Appendix. Due to the large number of parameters we only focus on the scalars μ_k, σ_k and ρ_k, as well as intensities in voxels v_1 and v_2 where the highest median posterior values are observed for the two types. Integrated intensities over the entire brain are also examined. Results indicate that our chains have converged to their stationary distribution. This is verified by the fact that posterior values from the 3 different runs overlap one with another for all the quantities that we plot. Finally it is worth noting that our chains show very good mixing properties since autocorrelation falls to low values after only a small number of iterations.

5.3 Results

In Figure 4 we plot median posterior intensities for emotion in several axial slices of the brain. Our results look qualitatively similar to the results obtained by Kober *et al.* (2008), Yue *et al.* (2012) and Kang *et al.* (2011) in their meta-analyses of emotion studies. The expected number of foci is 7.15 (95% CI [6.99,7.33]) and the regions mostly activated by emotions are the right and the left amygdala where the peak intensities in Figure 4 appear. Executive control processing generally recruits more regions of the brain and hence the median expected number of foci is 12.30 (95% CI [11.92,12.68]). The main effects are localised in the right and left cerebral cortex as can be seen in Figure 5.

Several quantities of interest can be obtained from our model, based on the properties of the spatial Poisson process. For example, one may calculate the probability of observing at least one focus in a set of voxels, e.g. a region of interest (ROI) or the entire brain. In Figure 7 we show the posterior distribution of $\mathbb{P}\left(N_{\mathbf{X}}\left(B\right) \geq 1\right)$, the probability of observing at least one focus in B, for several ROIs B. The division of the brain in ROIs is done according to the Harvard-Oxford atlas (Desikan *et al.*, 2006). A full brain analysis can be found in Section D of the Appendix. Note that this type of inference cannot be easily obtained from kernel-based methods such as MKDA (Wager *et al.*, 2007) or ALE (Eickhoff *et al.*, 2012) and therefore is a relative merit of our Bayesian point process model.

Comparison of the two types can be made as described in Section 4.2. The mean standardised posterior difference of the two intercepts is shown in Figure 6. Generally, we see that the two types have entirely distinct localisations and the main differences are observed where the main effects appear that is, the amygdala and regions of the cortex including front cortex, dorsolateral prefrontal cortex and parietal cortex.

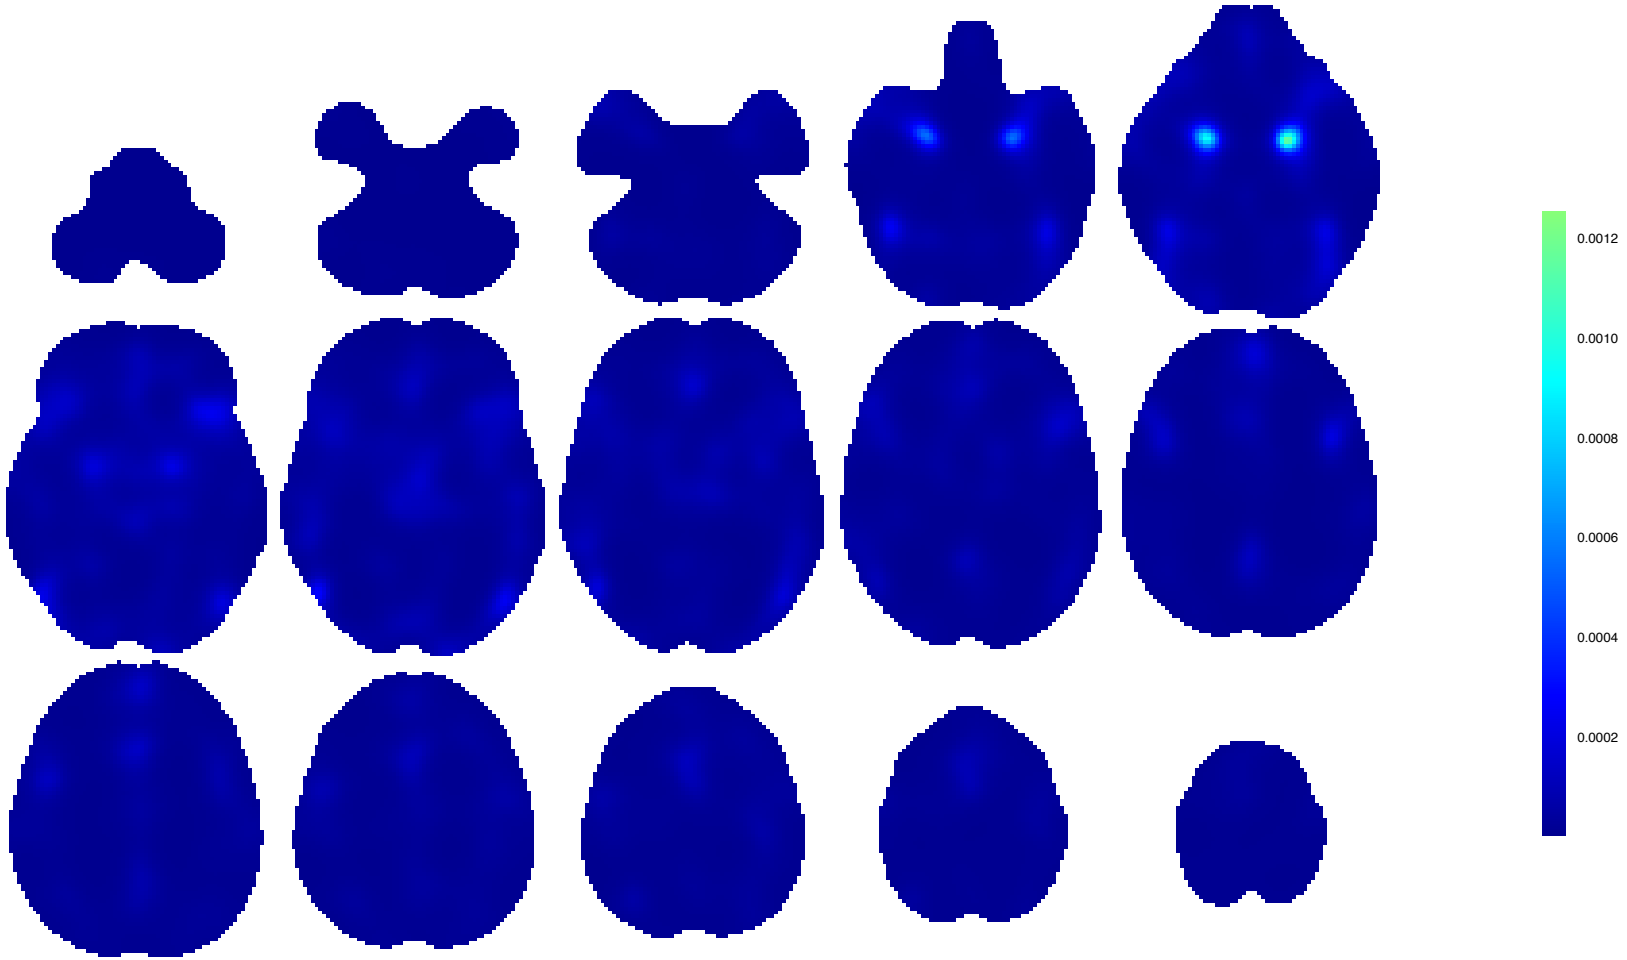

Figure 4: Emotion results. The figure presents the median posterior intensity at each voxel for several axial slices, as obtained from 1,000 independent draws from the posterior. Top row (left to right): $z = -50$, $z = -42$, $z = -32$, $z = -24$ and $z = -16$. Middle row (left to right): $z = -8$, $z = +2$, $z = +10$, $z = +18$ and $z = +28$. Bottom row (left to right): $z = +36$, $z = +44$, $z = +52$, $z = +62$ and $z = +70$.

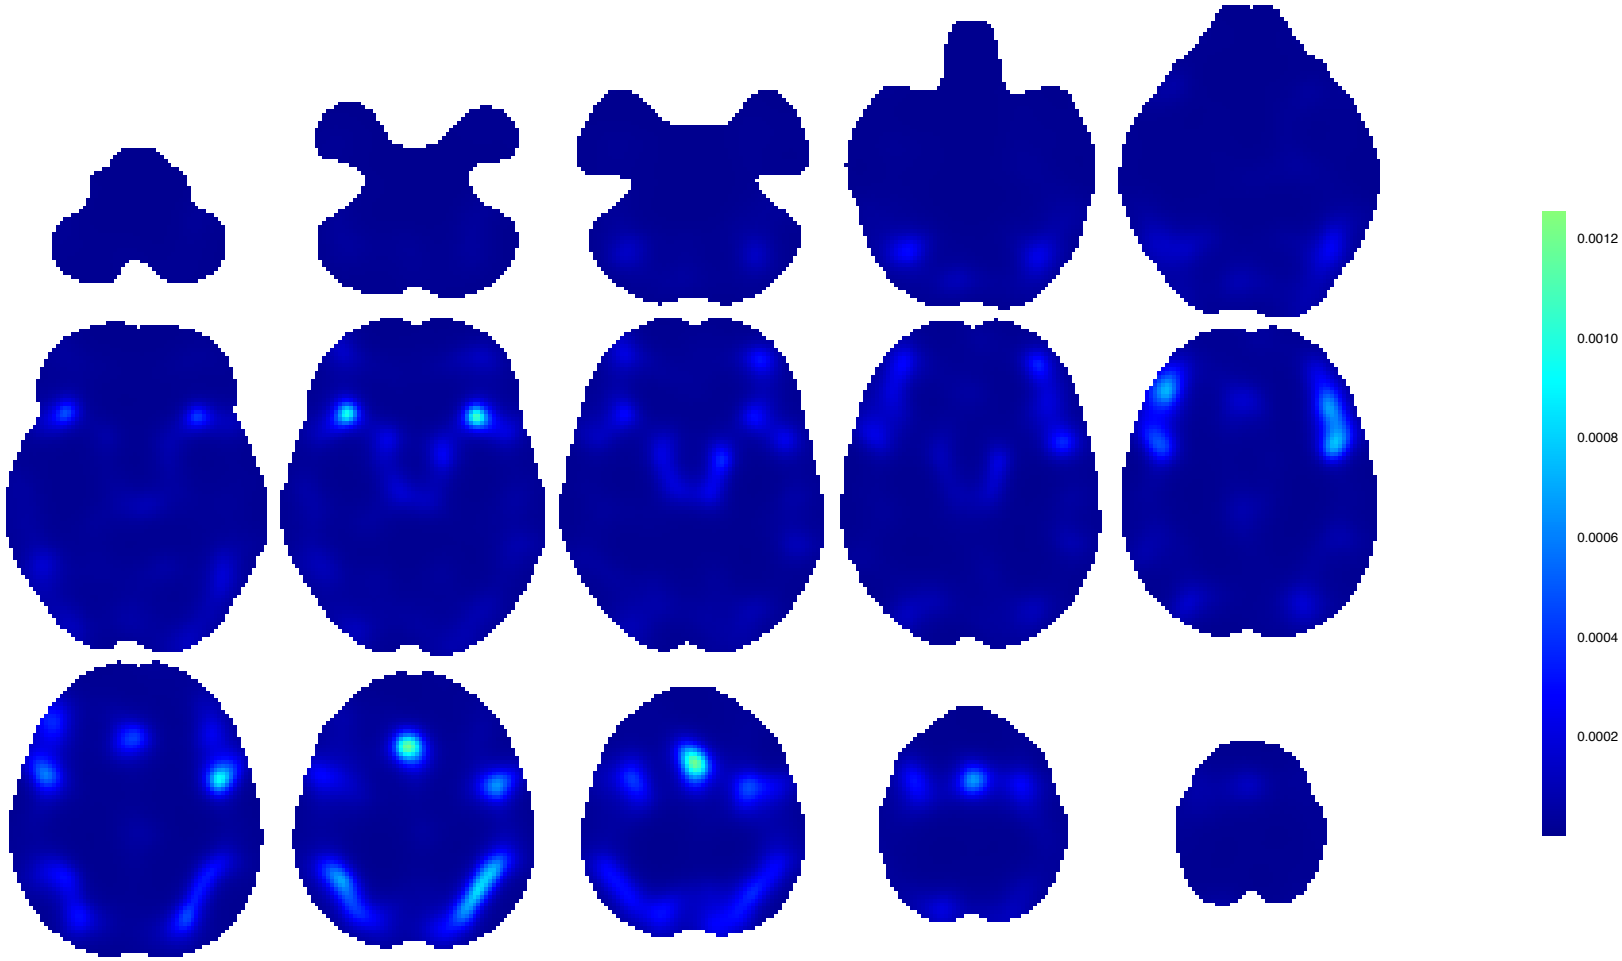

Figure 5: Executive control results. The figure presents the median posterior intensity at each voxel for several axial slices, as obtained from 1,000 independent draws from the posterior. Top row (left to right): $z = -50$, $z = -42$, $z = -32$, $z = -24$ and $z = -16$. Middle row (left to right): $z = -8$, $z = +2$, $z = +10$, $z = +18$ and $z = +28$. Bottom row (left to right): $z = +36$, $z = +44$, $z = +52$, $z = +62$ and $z = +70$.

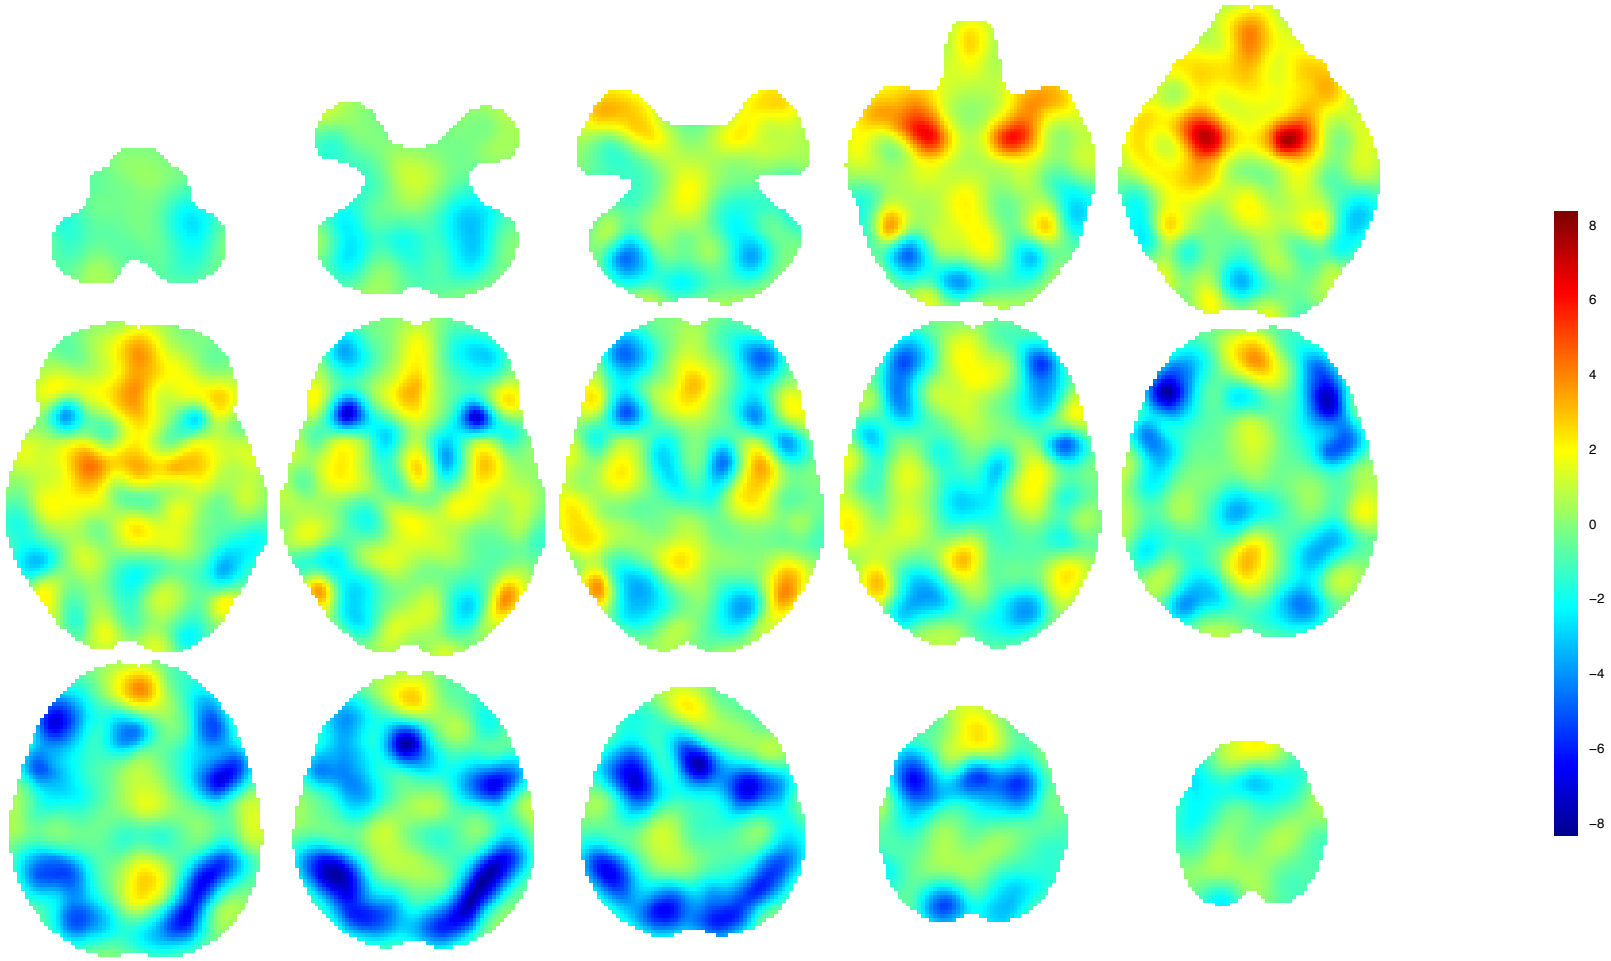

Figure 6: Comparison of the two types. The figure presents mean standardised posterior differences of the latent Gaussian processes at each voxel for several axial slices, as obtained from 1,000 independent draws from the posterior. Top row (left to right): $z = -50$, $z = -42$, $z = -32$, $z = -24$ and $z = -16$. Middle row (left to right): $z = -8$, $z = +2$, $z = +10$, $z = +18$ and $z = +28$. Bottom row (left to right): $z = +36$, $z = +44$, $z = +52$, $z = +62$ and $z = +70$.

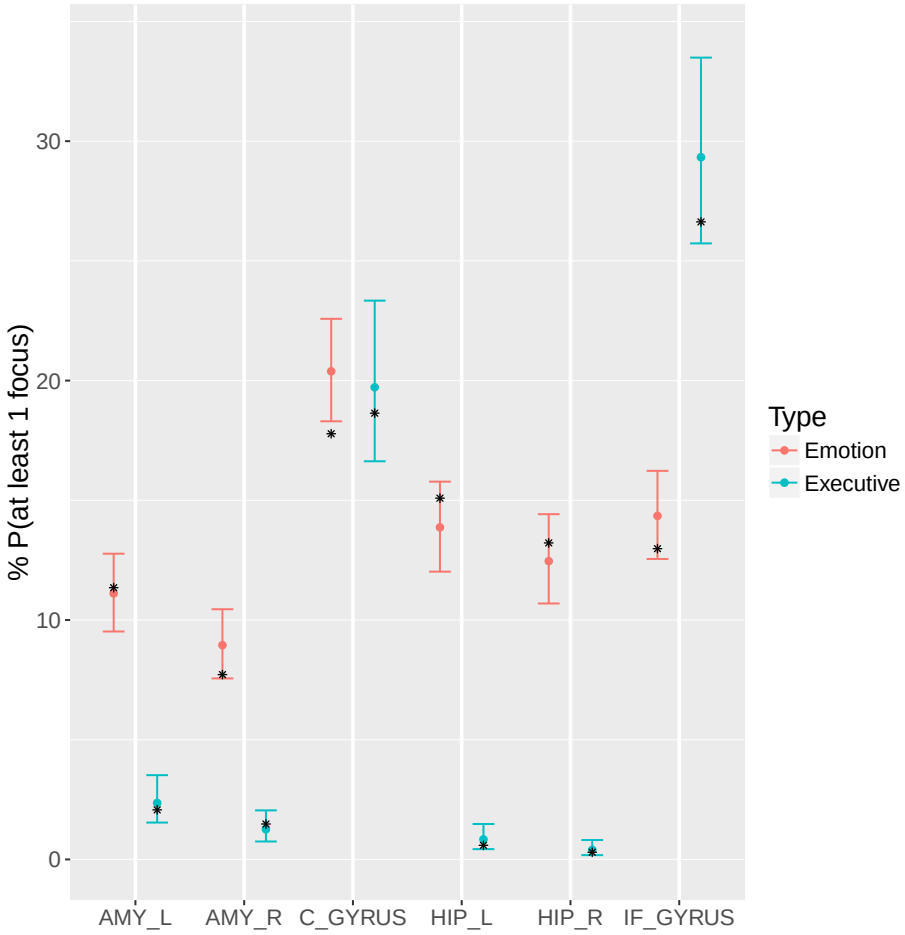

Figure 7: Posterior % probabilities of observing at least one focus for several ROIs. The points represent the median posterior values, red for emotion and blue for executive control. The errors represent the 95% posterior credible intervals. Black asterisks are the empirical values as obtained from the data. The ROIs are (left to right): left amygdala, right amygdala, anterior cingulate gyrus, left hippocampus, right hippocampus and inferior frontal gyrus pars opercularis.

5.4 Model assessment

Posterior predictive checks (Gelman *et al.*, 1996) can be found in Section E of the Appendix. In particular, we compare first and second order properties of observed data with samples obtained from the posterior predictive distribution. Results indicate some weakness in capturing the second order properties of the data but are overall satisfactory.

6 Discussion

In this work we have presented a new coordinate-based meta-analysis model, extension of the log-Gaussian Cox process model of Møller *et al.* (1998). To our knowledge, this is the first application of the LGCP with multiple realisations in a 3D problem. Note that even

though our application is focused on neuroimaging, the method is directly applicable to any multi-type point pattern problem.

The model has an appealing interpretation being a spatial GLM and several interesting inferences can be obtained based on the properties of the spatial Cox process. A significant advantage of our approach compared to existing methods is the inclusion of covariates in the analysis thus introducing the notion of meta-regression in CBMA. Another very interesting feature of the model is that it allows multi-type comparison directly from the posterior without having to run the model several times.

The main weakness of our approach is the large amount of computational effort required. Nevertheless, the proposed HMC algorithm exhibits quick convergence and good mixing properties and hence less samples from the posterior are required. Additionally, implementation of the method on a GPU vastly reduces computation time and thus makes the method applicable on big meta-analysis problems.

Application of the method on a meta-analysis of emotion and executive control studies has given valuable insights on the data. In particular, we have found that the main structures activated by emotion processing are the right and left amygdala, a finding which is consistent with previous studies on the topic. Executive control functions engage more regions compared to emotion and the main areas activated are the right and left cerebral cortex. Furthermore, comparison of the two types has allowed the detection of several regions in which they differ significantly.

There are several ways in which our work can be extended. A potential future direction is to study the conditions, such as sample size or minimum number of foci, under which it is possible to estimate several global or spatially varying effects using the LGCP. Such work can be of importance for practical implementations since it will provide some guidance regarding the complexity of meta-regression models that can be fit to a given dataset. One other option would be consider a random effects LGCP model where:

$$\log \lambda_i(\xi) = \sum_{k=0}^{K^*} \beta_k(\xi) z_{ik} + \sum_{k=K^*+1}^{K} \beta_k z_{ik} + \epsilon_i, \tag{13}$$

where ϵ_i are the random effect terms with $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$. This extension would account for further variability among the observed point patterns. Gradient expression are still available in closed form and hence inference for this model is still feasible by jointly updating the random effects parameters $\{\epsilon_i\}_{i=1}^I$ along with the remaining parameters of the LGCP.

Another possibility is to use some additional information about the foci such as p-values or T scores. These values can be attached as marks to the existing point patterns in order to improve estimation of the intensity function. Such an approach can enrich the inferences obtained from a meta-analysis by characterising the magnitude of activation in each region as opposed to the localisation of activations, which is the question that current methods address. Finally, it would be interesting to see how point pattern data can be

combined with IBMA data so that a meta-analysis can take advantage of all the available information on a topic. Radua *et al.* (2012) have already done some work for kernel-based methods but the problem still needs to be tackled for spatial models as well.

Acknowledgements

The authors would like to thank Derek Nee, John Jonides, David Hayes, Brian T. Gold, Georg Northoff and Raffael Kalisch for their help collecting and sharing the data. PS, TEN & TDJ were supported by NIH grant R01 NS075066-01A1; TEN was supported by a Wellcome Trust fellowship 100309/Z/12/Z.

References

Benes, V., Bodlák, K., Møller, J., and Waagepetersen, R. P. (2002). Bayesian analysis of log Gaussian Cox processes for disease mapping. Technical report, Depertment of Mathematical Sciences, Aalborg University.

Carp, J. (2012). The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage*, **63**(1), 289–300.

Christensen, O. F. and Waagepetersen, R. P. (2002). Bayesian prediction of spatial count data using generalized linear mixed models. *Biometrics*, **58**(2), 280–286.

Christensen, O. F., Roberts, G. O., and Sköld, M. (2006). Robust Markov chain Monte Carlo methods for spatial generalized linear mixed models. *Journal of Computational and Graphical Statistics*, **15**(1), 1–17.

Desikan, R. S., Sgonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., and Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, **31**(3), 968–980.

Dietrich, C. R. and Newsam, G. N. (1993). A fast and exact method for multidimensional Gaussian stochastic simulations. *Water Resources Research*, **29**(8), 2861–2869.

Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Physics Letters B*, **195**(2), 216–222.

Eickhoff, S. B., Bzdok, D., Laird, A. R., Kurth, F., and Fox, P. T. (2012). Activation likelihood estimation meta-analysis revisited. *NeuroImage*, **59**(3), 2349–2361.

Gelman, A., Meng, X.-L., and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, **6**(4), 733–807.

Girolami, M. and Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, **73**(2), 123–214.

Greenland, S. (1994). Invited commentary: a critical look at some popular meta-analytic methods. *American Journal of Epidemiology*, **140**(3), 290–296.

Hartung, J., Knapp, G., and Sinha, B. K. (2008). *Statistical Meta-Analysis with Applications*. John Wiley & Sons, Hoboken.

Hoffman, M. and Gelman, A. (2014). The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, **15**(1), 1593–1623.

Illian, J. B., Møller, J., and Waagepetersen, R. P. (2009). Hierarchical spatial point process analysis for a plant community with high biodiversity. *Environmental and Ecological Statistics*, **16**(3), 389–405.

Illian, J. B., Sørbye, S. H., and Rue, H. (2012a). A toolbox for fitting complex spatial point process models using integrated nested Laplace approximation (INLA). *The Annals of Applied Statistics*, **6**(4), 1499–1530.

Illian, J. B., Sørbye, S. H., Rue, H., and Hendrichsen, D. K. (2012b). Using INLA to fit a complex point process model with temporally varying effects – A case study. *Journal of Environmentl Statistics*, **3**(7), 1–25.

Jaakkola, T. and Jordan, M. (2000). Bayesian parameter estimation via variational methods. *Statistics and Computing*, **10**(1), 25–37.

Kang, J., Johnson, T. D., Nichols, T. E., and Wager, T. D. (2011). Meta analysis of functional neuroimaging data via Bayesian spatial point processes. *Journal of the American Statistical Association*, **106**(493), 124–134.

Kober, H., Barrett, L. F., Joseph, J., Bliss-Moreau, E., Lindquist, K., and Wager, T. D. (2008). Functional grouping and cortical and subcortical interactions in emotion: a meta-analysis of neuroimaging studies. *NeuroImage*, **42**(2), 998–1031.

Liang, S., Carlin, B. P., and Gelfand, A. E. (2009). Analysis of Minnesota colon and rectum cancer point patterns with spatial and nonspatial covariate information. *The Annals of Applied Statistics*, **3**(3), 943–962.

Lindquist, K. A., Wager, T. D., Kober, H., Bliss-Moreau, E., and Barrett, L. F. (2012). The brain basis of emotion: a meta-analytic review. *Behavioral and Brain Sciences*, **35**(3), 121–143.

Lindquist, M. A. (2008). The statistical analysis of fMRI data. *Statistical Science*, **23**(4), 439–464.

Marshall, T. and Roberts, G. (2012). An adaptive approach to Langevin MCMC. *Statistics and Computing*, **22**(5), 1041–1057.

Møller, J. and Waagepetersen, R. P. (2003). An introduction to simulation-based inference for spatial point processes. In J. Møller, editor, *Spatial Statistics and Computational Methods*, chapter 4, pages 143–198. Springer-Verlag.

Møller, J. and Waagepetersen, R. P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall/CRC, Boca Raton.

Møller, J. and Waagepetersen, R. P. (2007). Modern statistics for spatial point processes. *Scandinavian Journal of Statistics*, **34**(4), 643–684.

Møller, J., Syversveen, A. R., and Waagepetersen, R. P. (1998). Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, **25**(3), 451–482.

Murray, I., Adams, R. P., and MacKay, D. J. (2010). Elliptical slice sampling. *Journal of Machine Learning Research: Workshop and Conference Proceeding*, **9**(6), 541–548.

Neal, R. M. (1996). *Bayesian Learning for Neural Networks*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.

Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, and X. Meng, editors, *Handbook of Markov Chain Mote Carlo*, chapter 5, pages 113–162. Chapman & Hall/CRC.

Poldrack, R. A., Mumford, J. A., and Nichols, T. (2011). *Handbook of Functional MRI Data Analysis*. Cambridge University Press.

Radua, J. and Mataix-Cols, D. (2009). Voxel-wise meta-analysis of grey matter changes in obsessive-compulsive disorder. *The British Journal of Psychiatry : the Journal of Mental Science*, **195**(5), 393–402.

Radua, J., Mataix-Cols, D., Phillips, M. L., El-Hage, W., Kronhaus, D. M., Cardoner, N., and Surguladze, S. (2012). A new meta-analytic method for neuroimaging studies that combines reported peak coordinates and statistical parametric maps. *European Psychiatry*, **27**(8), 605–611.

Raemaekers, M., Vink, M., Zandbelt, B., van Wezel, R., Kahn, R., and Ramsey, N. (2007). Test-retest reliability of fMRI activation during prosaccades and antisaccades. *NeuroImage*, **36**(3), 532–542.

Raichle, M. E. (2003). Functional brain imaging and human brain function. *Journal of Neuroscience*, **23**(10), 3959–3962.

Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

Rottschy, C., Langner, R., Dogan, I., Reetz, K., Laird, A. R., Schulz, J. B., Fox, P. T., and Eickhoff, S. B. (2012). Modelling neural correlates of working memory: a coordinate-based meta-analysis. *NeuroImage*, **60**(1), 830–846.

Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.

Salimi-Khorshidi, G., Smith, S. M., Keltner, J. R., Wager, T. D., and Nichols, T. E. (2009). Meta-analysis of neuroimaging data: a comparison of image-based and coordinate-based pooling of studies. *NeuroImage*, **45**(3), 810–823.

Simpson, D., Illian, J., Lindgren, F., Sørbye, S., and Rue, H. (2016). Going off grid: computationally efficient inference for log-Gaussian Cox processes. *Biometrika*, **103**(1), 49–70.

Taylor, B. M. and Diggle, P. J. (2014). INLA or MCMC? A tutorial and comparative evaluation for spatial prediction in log-Gaussian Cox processes. *Journal of Statistical Computation and Simulation*, **84**(10), 2266–2284.

Turkeltaub, P. E., Eden, G. F., Jones, K. M., and Zeffiro, T. A. (2002). Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *NeuroImage*, **16**(3, Part A), 765–780.

Waagepetersen, R. P. (2004). Convergence of posteriors for discretized log Gaussian Cox processes. *Statistics and Probability Letters*, **66**(3), 229–235.

Wager, T. D., Jonides, J., and Reading, S. (2004). Neuroimaging studies of shifting attention: a meta-analysis. *NeuroImage*, **22**(4), 1679–1693.

Wager, T. D., Lindquist, M., and Kaplan, L. (2007). Meta-analysis of functional neuroimaging data: current and future directions. *Social Cognitive and Affective Neuroscience*, **2**(2), 150–158.

Wager, T. D., Lindquist, M. a., Nichols, T. E., Kober, H., and Van Snellenberg, J. X. (2009). Evaluating the consistency and specificity of neuroimaging data using meta-analysis. *NeuroImage*, **45**(Supplement 1), S210–S221.

Wood, A. T. A. and Chan, G. (1994). Simulation of stationary Gaussian processes in $[0, 1]^d$. *Journal of Computational and Graphical Statistics*, **3**(4), 409–432.

Yue, Y. R., Lindquist, M. A., and Loh, J. M. (2012). Meta-analysis of functional neuroimaging data using Bayesian nonparametric binary regression. *The Annals of Applied Statistics*, **6**(2), 697–718.

A Gradient expressions for the LGCP

The log-posterior, up to a normalising constant is given by:

$$\ell\left(\boldsymbol{\beta}, \boldsymbol{\sigma}, \boldsymbol{\rho}, \{\boldsymbol{\gamma}_k\}_{k=1}^{K^*} \mid \cdot\right) \propto \sum_{i=1}^{I}\left[-\sum_{j=1}^{V} A_{v_j}\lambda_i(v_j) + \sum_{j=1}^{V}\mathbf{1}_{v_j\in\mathbf{x}_i}\log\lambda_i(v_j)\right] + \log\text{priors}, \quad (14)$$

where $\boldsymbol{\beta} = [\beta_0 \equiv \mu_0, \ldots, \beta_{K^*} \equiv \mu_{K^*}, \beta_{K^*+1}, \ldots, \beta_K]^{\mathrm{T}}$, $\boldsymbol{\sigma} = [\sigma_0, \ldots, \sigma_{K^*}]^{\mathrm{T}}$, $\boldsymbol{\rho} = [\rho_0, \ldots, \rho_{K^*}]^{\mathrm{T}}$, $A_{v_j} = \mathbf{1}_{v_j\in\mathcal{B}}$ and the intensity function at each voxel v_j is defined as:

$$\lambda_i(v_j) = \exp\left(\sum_{k=0}^{K}\beta_k z_{ik}\right)\exp\left(\sum_{k=0}^{K^*}\sigma_k\left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_j z_{ik}\right) \quad (15)$$

We now calculate the derivatives with respect to the parameters of interest.

Partial derivatives with respect to β_l

We have that:

$$\begin{aligned}\frac{\partial\log\lambda_i(v_j)}{\partial\beta_l} &= \frac{\partial}{\partial\beta_l}\sum_{k=0}^{K}\beta_k z_{ik} + \frac{\partial}{\partial\beta_l}\sum_{k=0}^{K^*}\sigma_k\left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_j z_{ik}\\ &= z_{il}. \end{aligned} \quad (16)$$

As a result:

$$\begin{aligned}\frac{\partial\ell\left(\beta_l \mid \cdot\right)}{\partial\beta_l} &= -\sum_{i=1}^{I}\sum_{j=1}^{V}\left[A_{v_j}\frac{\partial}{\partial\beta_l}\lambda_i(v_j) - \mathbf{1}_{v_j\in\mathbf{x}_i}\frac{\partial}{\partial\beta_l}\log\lambda_i(v_j)\right] + \frac{\partial}{\partial\beta_l}\log\pi(\beta_l)\\ &= -\sum_{i=1}^{I}\sum_{j=1}^{V}\left[A_{v_j}\lambda_i(v_j)z_{il} - \mathbf{1}_{v_j\in\mathbf{x}_i}z_{il}\right] - \frac{\partial}{\partial\beta_l}\frac{\beta_l^2}{2\tau^2}\\ &= -\sum_{j=1}^{V}\sum_{i=1}^{I}\left[A_{v_j}\lambda_i(v_j)z_{il} + n_i z_{il}\right] - \frac{\beta_l}{\tau^2}, \end{aligned} \quad (17)$$

where n_i is the total number of foci in study i.

Partial derivatives with respect to σ_l

We have that:

$$\begin{aligned}\frac{\partial\log\lambda_i(v_j)}{\partial\sigma_l} &= \frac{\partial}{\partial\sigma_l}\sum_{k=0}^{K}\beta_k z_{ik} + \frac{\partial}{\partial\sigma_l}\sum_{k=0}^{K^*}\sigma_k\left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_j z_{ik}\\ &= \left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_j z_{ik}. \end{aligned} \quad (18)$$

Therefore:

$$\begin{aligned}
\frac{\partial \ell\left(\sigma_l \mid \cdot\right)}{\partial \sigma_l} &= -\sum_{i=1}^{I}\sum_{j=1}^{V}\left[A_{v_j}\frac{\partial}{\partial \sigma_l}\lambda_i(v_j) - \mathbf{1}_{v_j\in\mathbf{x}_i}\frac{\partial}{\partial \sigma_l}\log\lambda_i(v_j)\right] + \frac{\partial}{\partial \sigma_l}\log\pi(\sigma_l) \\
&= -\sum_{i=1}^{I}\sum_{j=1}^{V}\left[A_{vj}\lambda_i(v_j)\left(\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l\right)_j z_{il} - \mathbf{1}_{v_j\in\mathbf{x}_i}\left(\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l\right)_j z_{il}\right] - \frac{\partial}{\partial \sigma_l}\frac{\sigma_l^2}{2\tau^2} \\
&= -\sum_{j=1}^{V}\sum_{i=1}^{I}\left[\left(\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l\right)_j\left(A_{v_j}\lambda_i(v_j)z_{il} - \mathbf{1}_{v_j\in\mathbf{x}_i}z_{il}\right)\right] - \frac{\sigma_l}{\tau^2} \\
&= -\sum_{j=1}^{V}\left[\left(\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l\right)_j\sum_{i=1}^{I}\left[A_{v_j}\lambda_i(v_j)z_{il} - \mathbf{1}_{v_j\in\mathbf{x}_i}z_{il}\right]\right] - \frac{\sigma_l}{\tau^2}.
\end{aligned}$$

Partial derivatives with respect to ρ_l

Again:

$$\begin{aligned}
\frac{\partial \log\lambda_i(v_j)}{\partial \rho_l} &= \frac{\partial}{\partial \rho_l}\sum_{k=0}^{K^*}\beta_k z_{ik} + \frac{\partial}{\partial \rho_l}\sum_{k=0}^{K}\sigma_k\left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_j z_{ik} \\
&= \sigma_l\frac{\partial}{\partial \rho_l}\left(\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l\right)_j z_{il}. \qquad (19)
\end{aligned}$$

For ease of exposition we will complete the derivation for the one-dimensional case; however, similar arguments hold when $\mathcal{B}\subset\mathbb{R}^3$. Matrices $\mathbf{R}_l$ are circulant and so, the matrix-vector product $\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l$ can be found using the discrete Fourier transform as:

$$\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l = \mathbf{F}\boldsymbol{\Phi}_l^{1/2}\mathbf{F}^{\mathrm{H}}\boldsymbol{\gamma}_l, \qquad (20)$$

where $\boldsymbol{\Phi}_l$ are the diagonal matrices containing the eigenvalues of $\mathbf{R}_l$ and $\mathbf{F}$ is the matrix of eigenvectors. In Equation (20), the only term depending on ρ_l is $\boldsymbol{\Phi}_l$ and, hence:

$$\frac{\partial}{\partial \rho_l}\mathbf{R}_l^{1/2}\boldsymbol{\gamma}_l = \mathbf{F}\frac{\partial}{\partial \rho_l}\boldsymbol{\Phi}_l^{1/2}\mathbf{F}^{\mathrm{H}}\boldsymbol{\gamma}_l \qquad (21)$$

We know that $\boldsymbol{\Phi}_l = \mathrm{diag}\left\{\phi_{l_0},\ldots,\phi_{l_{V-1}}\right\}$, where for $k=0,\ldots,V-1$ we have that:

$$\phi_{l_k} = \sum_{j=0}^{V-1}\exp\left(-\rho_l||v_0,v_j||^{\delta_l}\right)\exp\left(-\frac{2\pi\iota kj}{V}\right), \qquad (22)$$

where ι is the imaginary unit. Now it is straightforward to see that for $k = 0, \ldots, V-1$:

$$\begin{aligned}
\frac{\partial}{\partial \rho_l} \phi_{l_k}^{1/2} &= \frac{\partial}{\partial \rho_l} \sqrt{\sum_{j=0}^{V-1} \exp\left(-\rho_l ||v_0, v_j||^{\delta_l}\right) \exp\left(-\frac{2\pi\iota kj}{V}\right)} \\
&= \frac{\frac{\partial}{\partial \rho_l} \sum_{j=0}^{V-1} \exp\left(-\rho_l ||v_0, v_j||^{\delta_l}\right) \exp\left(-\frac{2\pi\iota kj}{V}\right)}{2\sqrt{\sum_{j=0}^{V-1} \exp\left(-\rho_l ||v_0, v_j||^{\delta_l}\right) \exp\left(-\frac{2\pi\iota kj}{V}\right)}} \\
&= \frac{-\sum_{j=1}^{V-1} d\left(v_0, v_j\right)^{\delta_l} \exp\left(-\rho_l ||v_0, v_j||^{\delta_l}\right) \exp\left(-\frac{2\pi\iota kj}{V}\right)}{2\phi_{l_k}^{1/2}} \\
&= -\frac{1}{2} \frac{\psi_{l_k}}{\phi_{l_k}^{1/2}}, \qquad (23)
\end{aligned}$$

where ψ_{l_k} can be viewed as the k-th eigenvalue of the of a circulant matrix $\mathbf{S}_l$ with base $\mathbf{s}_l = \left[||v_0, v_0||^{\delta_l} \exp\left(-\rho_l ||v_0, v_0||^{\delta_l}\right), \ldots, ||v_0, v_{V-1}||^{\delta_l} \exp\left(-\rho_l ||v_0, v_{V-1}||^{\delta_l}\right)\right]$ and $\mathbf{S}_l = \mathbf{F}\boldsymbol{\Psi}_l\mathbf{F}^{\mathrm{H}}$, $\boldsymbol{\Psi}_l = \mathrm{diag}\left\{\psi_{l_0}, \ldots, \psi_{l_{V-1}}\right\}$. Overall we see that:

$$\begin{aligned}
\frac{\partial}{\partial \rho_l} \mathbf{R}_l^{1/2} \boldsymbol{\gamma}_l &= \mathbf{F} \frac{\partial}{\partial \rho_l} \boldsymbol{\Phi}_l^{1/2} \mathbf{F}^{\mathrm{H}} \boldsymbol{\gamma}_l \\
&= -\frac{1}{2} \mathbf{F} \left[\boldsymbol{\Psi}_l \oslash \boldsymbol{\Phi}^{1/2}\right] \mathbf{F}^{\mathrm{H}} \boldsymbol{\gamma}_l \\
&= -\frac{1}{2} \mathbf{Q}_l \boldsymbol{\gamma}_l, \qquad (24)
\end{aligned}$$

where $\oslash$ stands for element wise division. Combining Equations (19) and (24), we find that:

$$\frac{\partial \log \lambda_i(v_j)}{\partial \rho_l} = -\frac{1}{2} \sigma_l \left(\mathbf{Q}_l \boldsymbol{\gamma}_l\right)_j z_{il}. \qquad (25)$$

So:

$$\begin{aligned}
\frac{\partial \ell\left(\rho_l \mid \cdot\right)}{\partial \rho_l} &= -\sum_{i=1}^{I} \sum_{j=1}^{V} \left[A_{v_j} \frac{\partial}{\partial \rho_l} \lambda_i(v_j) - \mathbf{1}_{v_j \in \mathbf{x}_i} \frac{\partial}{\partial \rho_l} \log \lambda_i(v_j)\right] + \frac{\partial}{\partial \rho_l} \log \pi(\rho_l) \\
&= -\sum_{i=1}^{I} \sum_{j=1}^{V} \left[A_{vj} \lambda_i(v_j) \left(-\frac{1}{2}\right) \sigma_l \left(\mathbf{Q}_l \boldsymbol{\gamma}_l\right)_j z_{il} - \mathbf{1}_{v_j \in \mathbf{x}_i} \left(-\frac{1}{2}\right) \sigma_l \left(\mathbf{Q}_l \boldsymbol{\gamma}_l\right)_j z_{il}\right] \\
&\quad -\frac{\partial}{\partial \rho_l} 1 \\
&= \frac{\sigma_l}{2} \sum_{j=1}^{V} \sum_{i=1}^{I} \left[\left(\mathbf{Q}_l \boldsymbol{\gamma}_l\right)_j \left(A_{v_j} \lambda_i(v_j) z_{il} - \mathbf{1}_{v_j \in \mathbf{x}_i} z_{il}\right)\right] \\
&= \frac{\sigma_l}{2} \sum_{j=1}^{V} \left[\left(\mathbf{Q}_l \boldsymbol{\gamma}_l\right)_j \sum_{i=1}^{I} \left[A_{v_j} \lambda_i(v_j) z_{il} - \mathbf{1}_{v_j \in \mathbf{x}_i} z_{il}\right]\right].
\end{aligned}$$

Partial derivatives with respect to $\boldsymbol{\gamma}_l$

Finally:

$$\begin{aligned}\frac{\partial \log \lambda_i(v_j)}{\partial \boldsymbol{\gamma}_l} &= \frac{\partial}{\partial \boldsymbol{\gamma}_l}\sum_{k=0}^{K^*}\beta_k z_{ik} + \frac{\partial}{\partial \boldsymbol{\gamma}_l}\sum_{k=0}^{K}\sigma_k\left(\mathbf{R}_k^{1/2}\boldsymbol{\gamma}_k\right)_j z_{ik} \\ &= \sigma_l \mathbf{r}_{l_j} z_{il}, \end{aligned} \tag{26}$$

where $\mathbf{r}_{l_j}$ is the j-th row of the matrix $\mathbf{R}_l^{1/2}$. Now we can see that:

$$\begin{aligned}\frac{\partial \ell\left(\boldsymbol{\gamma}_l \mid \cdot\right)}{\partial \boldsymbol{\gamma}_l} &= -\sum_{i=1}^{I}\sum_{j=1}^{V}\left[A_{v_j}\frac{\partial}{\partial \boldsymbol{\gamma}_l}\lambda_i(v_j) - \mathbf{1}_{v_j\in\mathbf{x}_i}\frac{\partial}{\partial \boldsymbol{\gamma}_l}\log\lambda_i(v_j)\right] + \frac{\partial}{\partial \boldsymbol{\gamma}_l}\log\pi(\boldsymbol{\gamma}_l) \\ &= -\sum_{i=1}^{I}\sum_{j=1}^{V}\left[A_{v_j}\lambda_i(v_j)\sigma_l\mathbf{r}_{l_j}z_{il} - \mathbf{1}_{v_j\in\mathbf{x}_i}\sigma_l\mathbf{r}_{l_j}z_{il}\right] - \frac{\partial}{\partial \boldsymbol{\gamma}_l}\frac{\boldsymbol{\gamma}_l^{\mathrm{T}}\boldsymbol{\gamma}_l}{2} \\ &= -\sigma_l\sum_{j=1}^{V}\left[\mathbf{r}_{l_j}\sum_{i=1}^{I}\left[A_{v_j}\lambda_i(v_j)z_{il} - \mathbf{1}_{v_j\in\mathbf{x}_i}z_{il}\right]\right] - \boldsymbol{\gamma}_l \\ &= -\sigma_l\sum_{j=1}^{V}\left[\mathbf{r}_{l_j}\mathbf{c}_{l_j}\right] - \boldsymbol{\gamma}_l \\ &= -\sigma_l{\mathbf{R}_l^{1/2}}^{\mathrm{T}}\mathbf{c}_l - \boldsymbol{\gamma}_l \\ &= -\sigma_l\mathbf{R}_l^{1/2}\mathbf{c}_l - \boldsymbol{\gamma}_l, \end{aligned} \tag{27}$$

since $\mathbf{R}$ is a nested block circulant matrix, where $\mathbf{c}_l$ are V-vectors with elements $\mathbf{c}_{l_j} = \sum_{i=1}^{I}\left[A_{v_j}\lambda_i(v_j)z_{il} - \mathbf{1}_{v_j\in\mathbf{x_i}}z_{il}\right]$.

B LGCP simulation setup I traceplots

In this section we provide traceplots for the scalar parameters of the LGCP model, as fit to the simulated data of Section 4.1. Trace plots for the parameters σ_1, σ_2, ρ_1 and ρ_2 can be found in Figure 8, whereas trace plots for μ_1, μ_2, β_3 and β_4 can be found in Figure 9. The red lines indicate the true parameter values.

C Real data analysis diagnostics

This section contains convergence diagnostics for the real data analysis of Section 5. All plots are obtained from 3 different runs of HMC, from which we discard the first 5,000 draws as a burn-in. Parameters σ_k, ρ_k and μ_k, $k = 1, 2$, are shown in Figures 10, 11 and 12 respectively. Figure 13 shows draws from the posterior expected number of points per type. In Figure 14, we plot posterior intensities for voxel v_1 where the highest median posterior

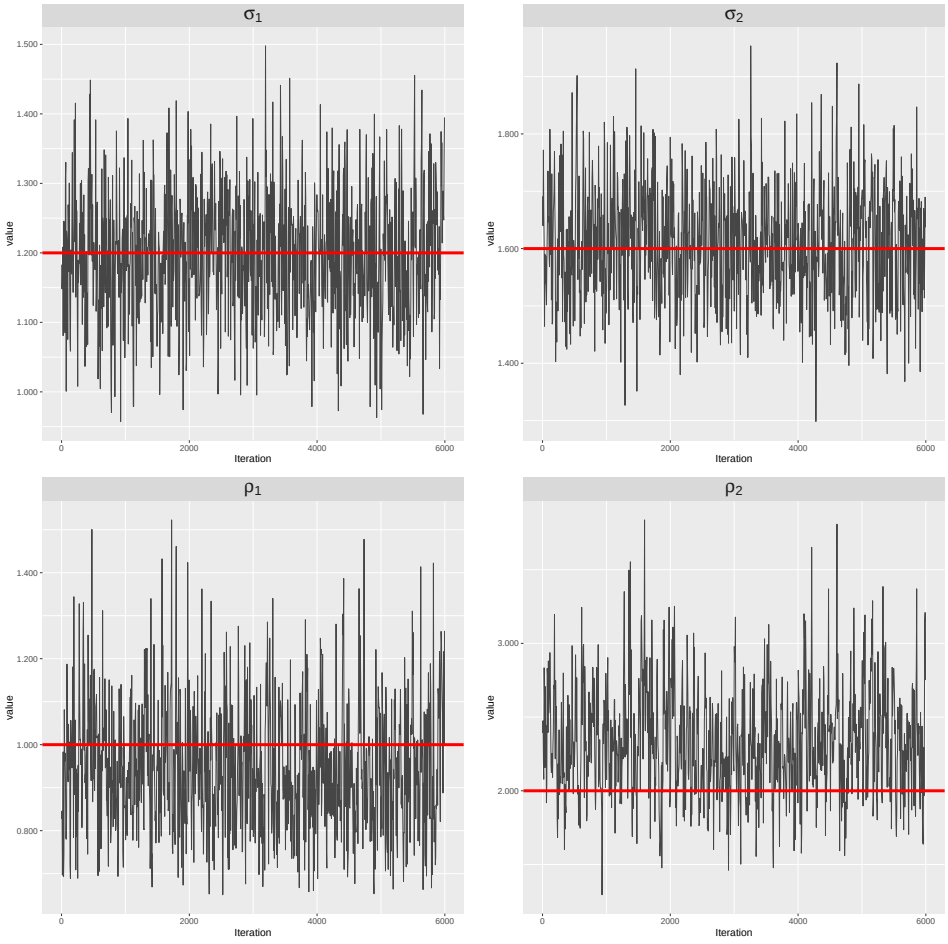

Figure 8: Posterior traceplots for the scalar parameters of the LGCP model used to fit the data of Section 4.1. Top row: standard deviations. Bottom row: correlation decay parameters ($\times 100$). The true values are indicated by the solid red lines.

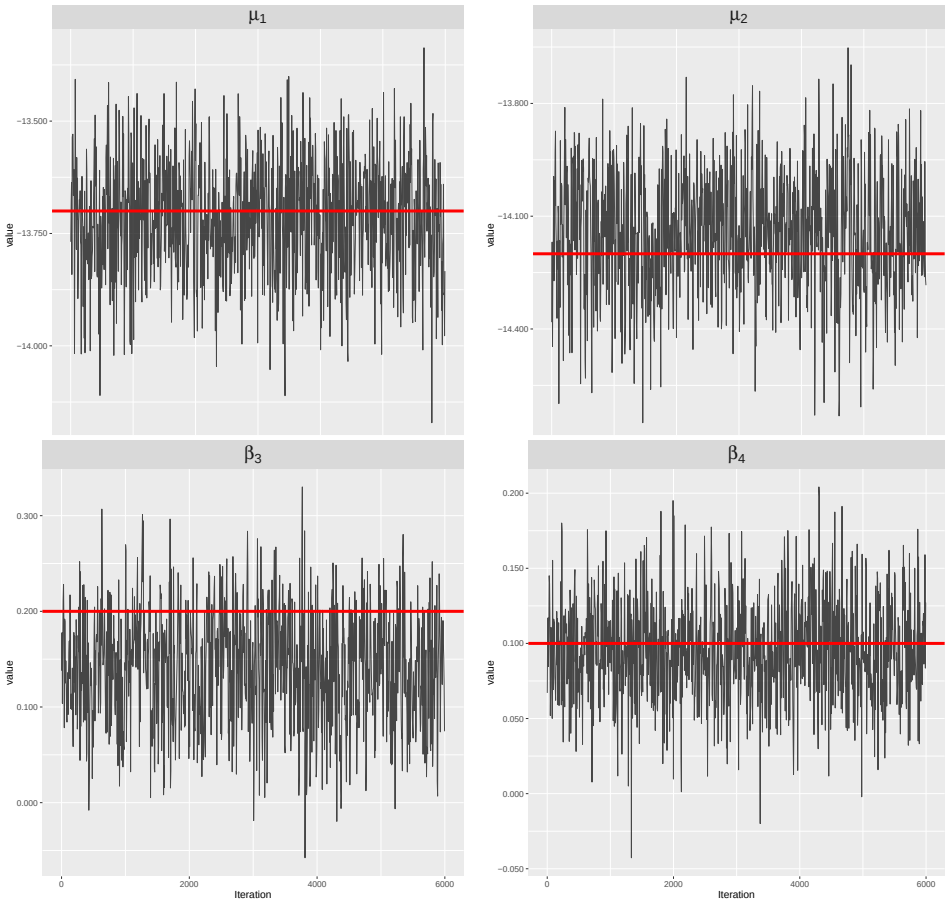

Figure 9: Posterior traceplots for the scalar parameters of the LGCP model used to fit the data of Section 4.1. Top row: overall latent process means. Bottom row: regression coefficients for covariates z_3 and z_4. The true values are indicated by the solid red lines.

intensity for emotions is observed. Finally, in Figure 15 we do the same for voxel v_2 which corresponds to the highest median posterior intensity for executive control studies.

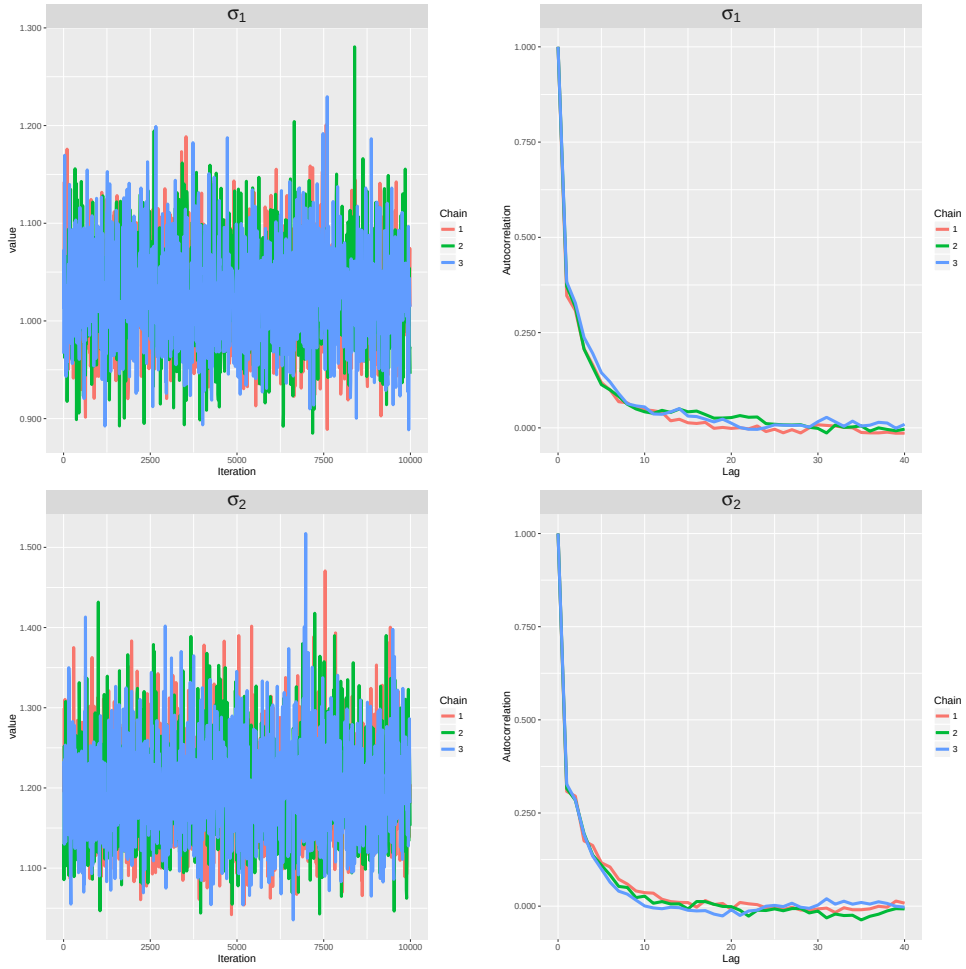

Figure 10: Posterior traceplots (left) along with autocorrelation plots (right) for the marginal standard deviations σ_k, as obtained from application of the LGCP meta-analysis model to the real data of Section 5. Top row corresponds to emotion while bottom row to executive control.

D Full brain analysis

In this section we present full brain results for the application of Section 5. Table 2 show posterior summaries for emotions whereas Table 3 shows posterior summaries for executive control. ROIs are obtained from the Harvard-Oxford atlas (Desikan *et al.*, 2006). The ROI volumes presented are in voxels. All quantities are based on 1,000 HMC samples, obtained after thinning the original chain every 10 iterations.

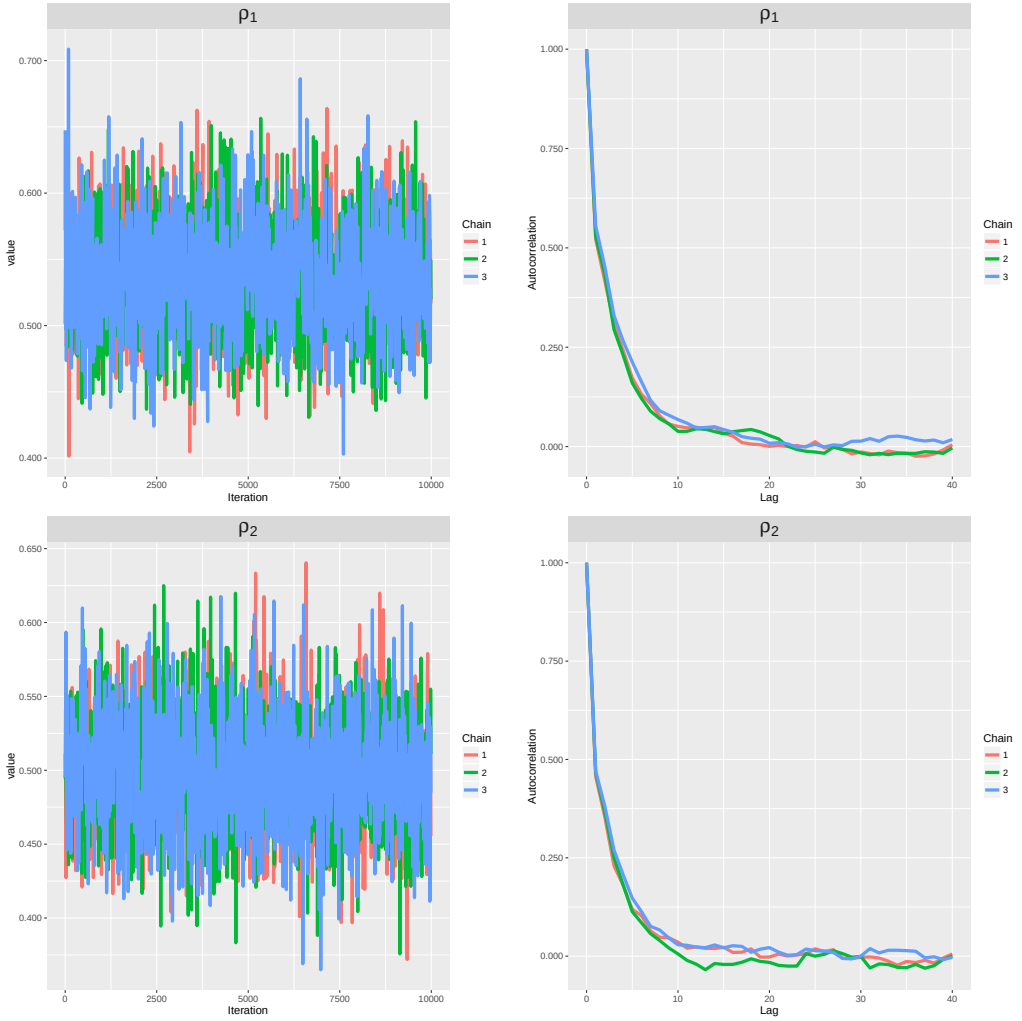

Figure 11: Posterior traceplots (left) along with autocorrelation plots (right) for the correlation decay parameters ρ_k, as obtained from application of the LGCP meta-analysis model to the real data of Section 5. Top row corresponds to emotion while bottom row to executive control. The actual values are multiplied by 100.

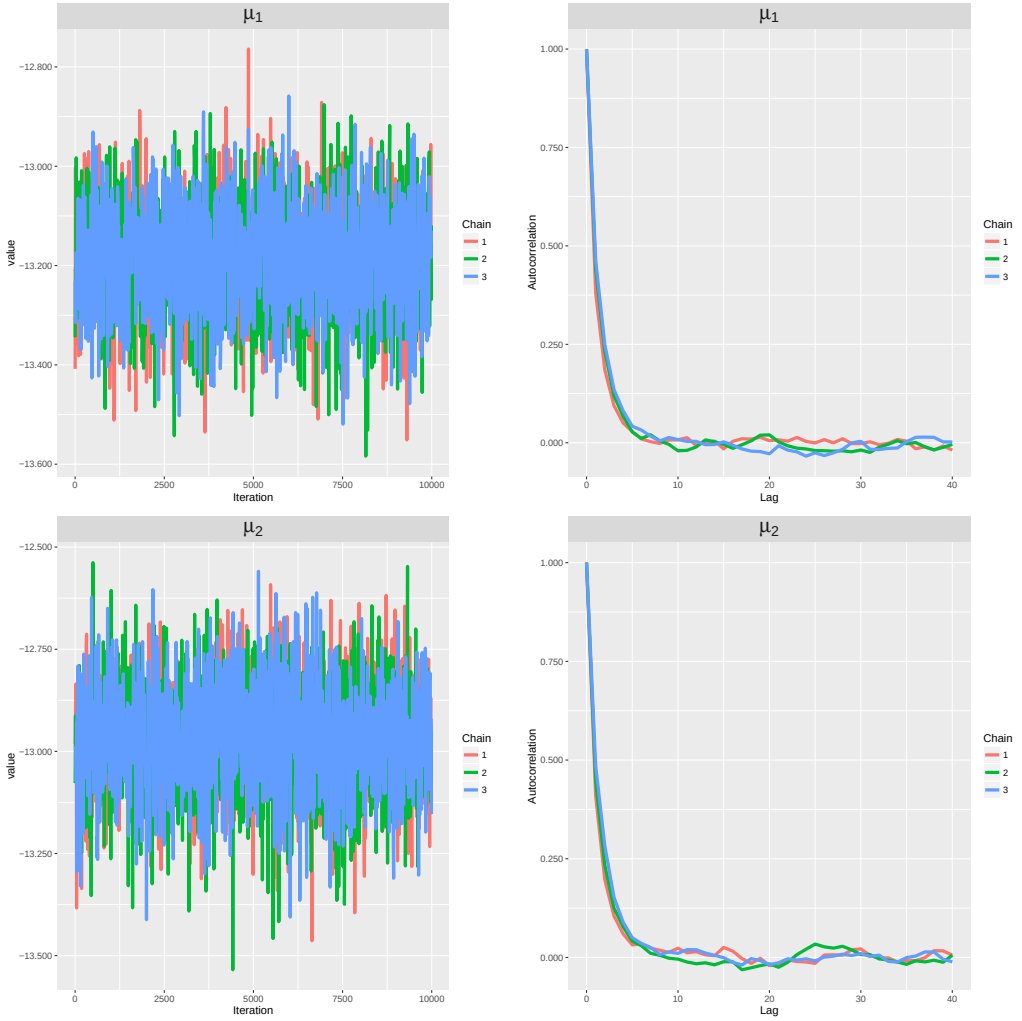

Figure 12: Posterior traceplots (left) along with autocorrelation plots (right) for the overall latent process means μ_k, as obtained from application of the LGCP meta-analysis model to the real data of Section 5. Top row corresponds to emotion while bottom row to executive control.

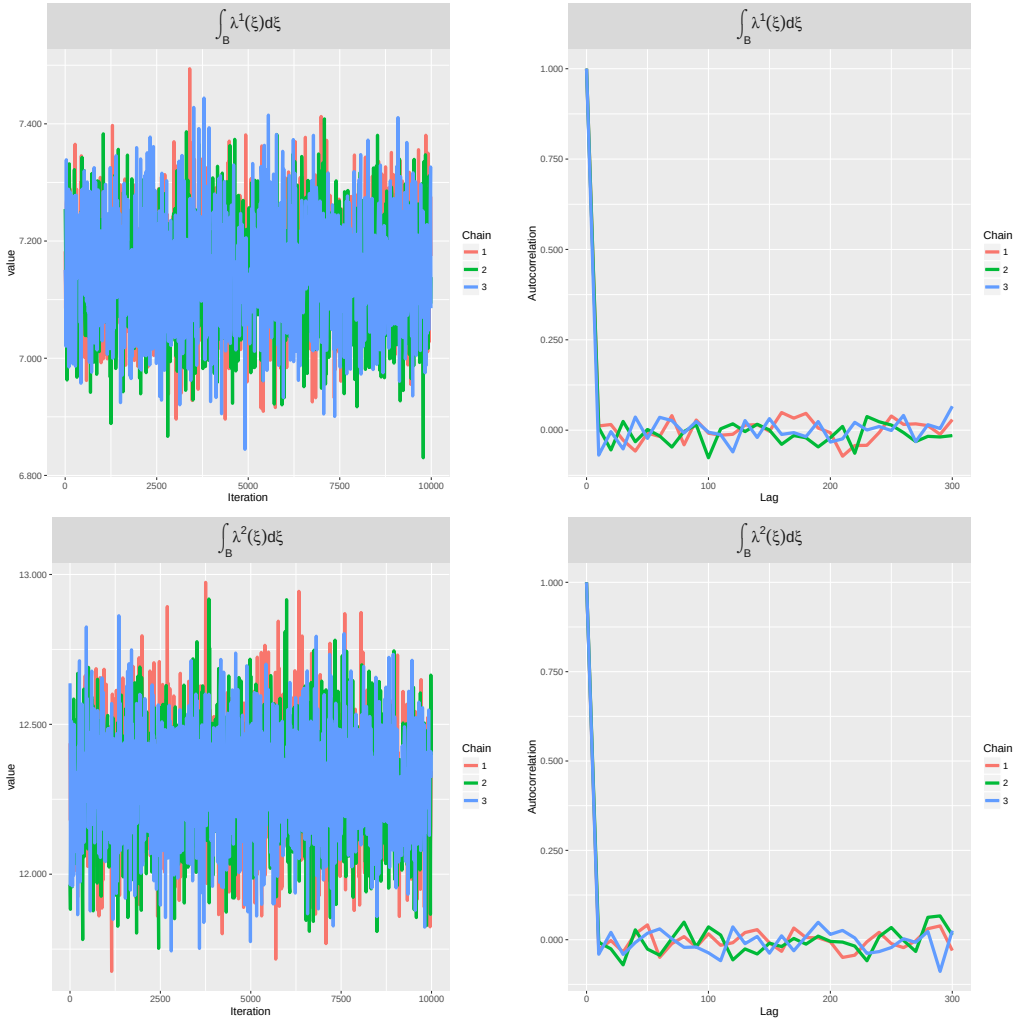

Figure 13: Posterior traceplots (left) along with autocorrelation plots (right) for the integrated intensities over the brain, as obtained from application of the LGCP meta-analysis model to the real data of Section 5. Top row corresponds to emotion while bottom row to executive control. Results are obtained after applying a thinning factor of 10 to the original draws.

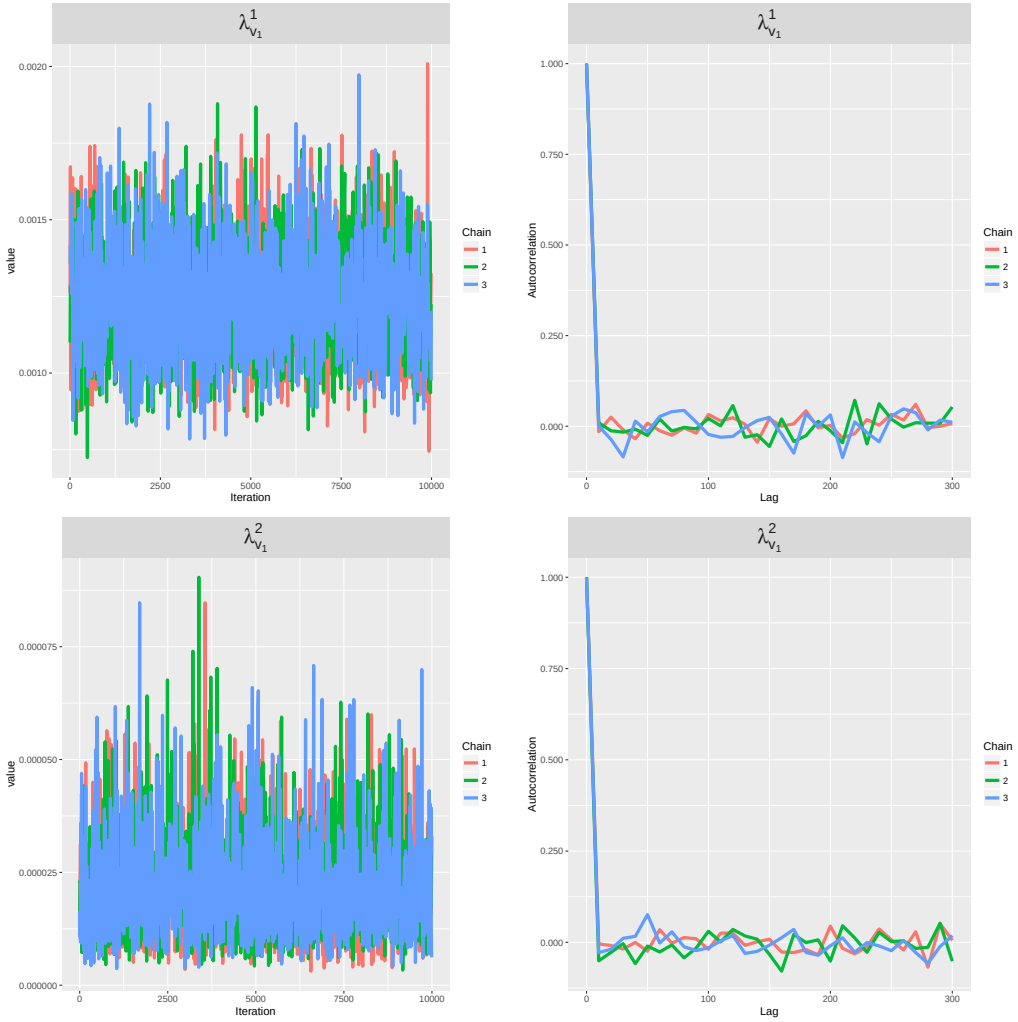

Figure 14: Posterior traceplots (left) along with autocorrelation plots (right) for intensities at voxel v_1 where the highest median posterior intensity was observed for emotions, as obtained from application of the LGCP meta-analysis model to the real data of Section 5. Top row corresponds to emotion while bottom row to executive control. Results are obtained after applying a thinning factor of 10 to the original draws.

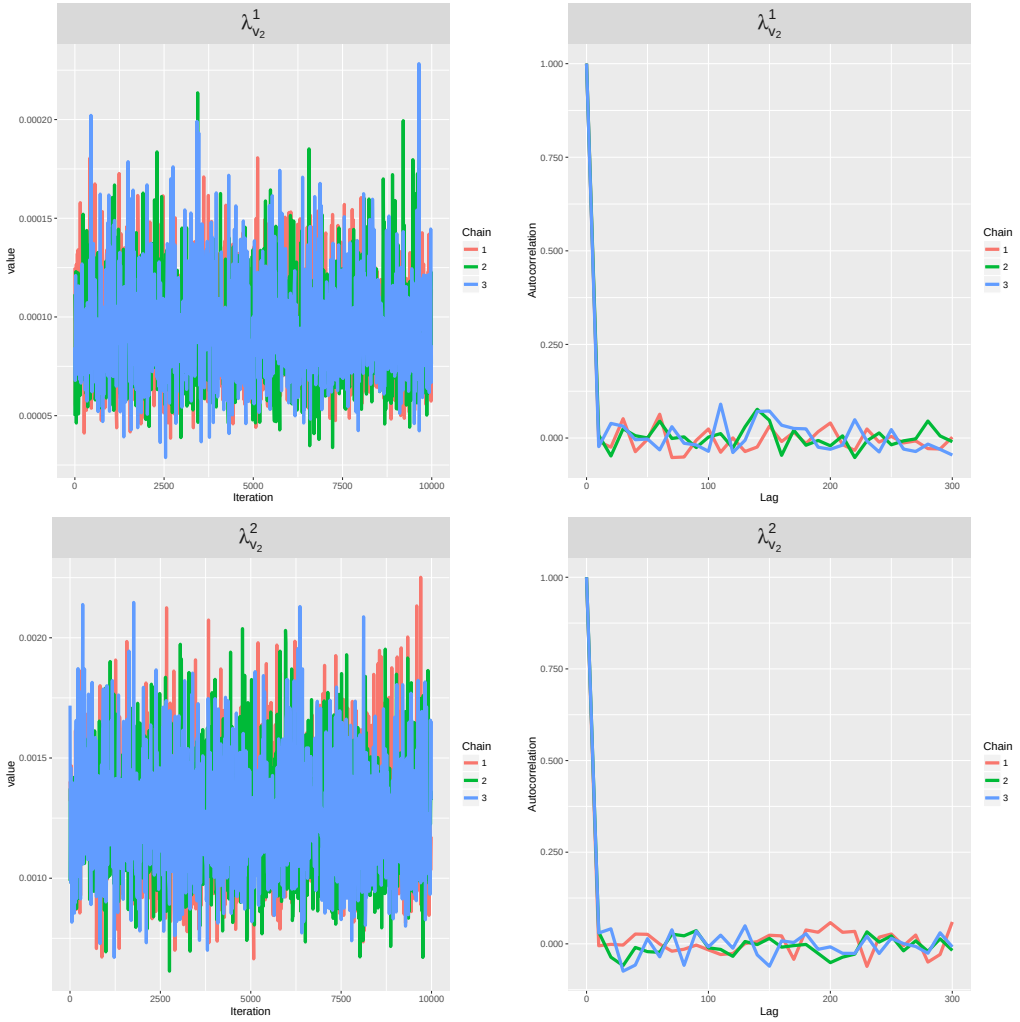

Figure 15: Posterior traceplots (left) along with autocorrelation plots (right) for intensities at voxel v_2 where the highest median posterior intensity was observed for executive control, as obtained from application of the LGCP meta-analysis model to the real data of Section 5. Top row corresponds to emotion while bottom row to executive control. Results are obtained after applying a thinning factor of 10 to the original draws.

Table 2: Meta-analysis results for emotions. Table presents posterior summaries for the probability of at least one activation in a given ROI (rounded to 2 decimal points), as well as the expected number points (rounded to 3 decimal points). The empirical quantities are also presented. Note the abbreviations: (L) for left, (R) for right, (A) for anterior, (P) for posterior, (S) for superior and (I) for inferior. All results are based on 1,000 HMC draws which are obtained after thinning the initial run every 10 iterations.

		$\% \, \mathbb{P}\left(\mathbf{N_X}\left(\mathbf{B}\right) \geq \mathbf{1}\right)$				$\int_{\mathbf{B}} \lambda\left(\xi\right) \mathbf{d}\xi$			
ROI	**VOL.**	$\mathbf{p_{0.5}}$	$\mathbf{p_{0.025}}$	$\mathbf{p_{0.975}}$	**DATA**	$\mathbf{p_{0.5}}$	$\mathbf{p_{0.025}}$	$\mathbf{p_{0.975}}$	**DATA**
Lateral ventrical (L)	1289	2.51	1.96	3.17	2.34	0.025	0.020	0.032	0.028
Lateral ventrical (R)	1019	1.99	1.54	2.50	1.64	0.020	0.016	0.025	0.020
Thalamus (L)	1591	10.48	8.85	12.40	10.29	0.111	0.093	0.132	0.120
Thalamus (R)	1398	7.97	6.53	9.61	8.65	0.083	0.068	0.101	0.092
Caudate (L)	572	2.04	1.53	2.74	2.11	0.021	0.015	0.028	0.022
Caudate (R)	515	1.70	1.23	2.29	1.99	0.017	0.012	0.023	0.022
Putamen (L)	923	4.49	3.48	5.64	4.91	0.046	0.035	0.058	0.051
Putamen (R)	800	3.58	2.85	4.52	3.86	0.036	0.029	0.046	0.042
Pallidum (L)	312	1.42	1.01	1.95	0.70	0.014	0.010	0.020	0.007
Pallidum (R)	266	1.61	1.19	2.21	1.52	0.016	0.012	0.022	0.016
Brain stem	7413	11.56	9.96	13.34	9.24	0.123	0.105	0.143	0.120
Hippocampus (L)	921	11.11	9.52	12.77	11.35	0.118	0.100	0.137	0.123
Hippocampus (R)	772	8.95	7.56	10.45	7.72	0.094	0.079	0.110	0.087
Amygdala (L)	390	13.87	12.02	15.78	15.09	0.149	0.128	0.172	0.159
Amygdala (R)	399	12.46	10.69	14.42	13.22	0.133	0.113	0.156	0.138
Accumbens (L)	111	0.60	0.39	0.90	1.05	0.006	0.004	0.009	0.011
Accumbens (R)	86	0.41	0.27	0.62	0.47	0.004	0.003	0.006	0.005
Frontal pole	15188	29.16	26.87	31.76	23.04	0.345	0.313	0.382	0.331
Insular cortex	3613	19.50	17.63	21.73	15.09	0.217	0.194	0.245	0.226

Superior frontal gyrus	7295	18.13	16.19	20.26	14.97	0.200	0.177	0.226	0.204
Middle frontal gyrus	6722	12.46	10.85	14.34	9.47	0.133	0.115	0.155	0.113
Inferior frontal gyrus, pars triangularis	1615	10.05	8.59	11.62	9.94	0.106	0.090	0.124	0.109
Inferior frontal gyrus, pars opercularis	1877	14.35	12.55	16.23	12.98	0.155	0.134	0.177	0.172
Precentral gyrus	11055	25.76	23.44	28.29	20.70	0.298	0.267	0.333	0.322
Temporal pole	3341	17.78	15.73	20.02	14.97	0.196	0.171	0.223	0.192
Superior temporal gyrus (A)	622	2.20	1.69	2.87	2.11	0.022	0.017	0.029	0.025
Superior temporal gyrus (P)	2093	8.83	7.41	10.17	7.25	0.092	0.077	0.107	0.090
Middle temporal gyrus (A)	762	2.73	2.05	3.63	2.69	0.028	0.021	0.037	0.028
Middle temporal gyrus (P)	2794	8.37	6.98	9.83	7.60	0.087	0.072	0.103	0.094
Middle temporooccipital temporal gyrus	2470	10.48	8.93	12.17	9.47	0.111	0.094	0.130	0.119
Inferior temporal gyrus (A)	789	1.18	0.80	1.69	1.05	0.012	0.008	0.017	0.012
Inferior temporal gyrus (P)	2687	3.08	2.33	3.99	1.99	0.031	0.024	0.041	0.020
Inferior temporoocipital temporal gyrus	1883	5.06	4.11	6.20	3.86	0.052	0.042	0.064	0.042
Postcentral gyrus	7553	12.30	10.67	14.09	8.19	0.131	0.113	0.152	0.112
Superior parietal lobule	3626	4.91	3.86	6.08	4.21	0.050	0.039	0.063	0.047
Supramarginal gyrus, (A)	1951	4.30	3.35	5.49	4.21	0.044	0.034	0.056	0.051
Supramarginal gyrus, (P)	3017	6.10	5.04	7.31	5.50	0.063	0.052	0.076	0.063
Angular gyrus	2944	5.98	4.86	7.25	5.15	0.062	0.050	0.075	0.058
Lateral occipital cortex (S)	9535	17.99	15.92	20.27	13.45	0.198	0.173	0.226	0.198
Lateral occipital cortex (I)	5004	30.57	28.02	33.29	22.46	0.365	0.329	0.405	0.382
Intracalcarine cortex	2211	4.89	3.92	6.09	4.44	0.050	0.040	0.063	0.048
Frontal medial cortex	1473	6.35	5.18	7.72	5.50	0.066	0.053	0.080	0.061
Juxtapositional lobule cortex	2193	7.97	6.56	9.59	8.30	0.083	0.068	0.101	0.091
Subcallosal cortex	2150	7.21	5.81	8.81	5.96	0.075	0.060	0.092	0.078

Paracingulate gyrus	4095	18.71	16.74	20.88	15.67	0.207	0.183	0.234	0.202
Cingulate gyrus (A)	4144	20.39	18.30	22.58	17.78	0.228	0.202	0.256	0.237
Cingulate gyrus (P)	4668	10.22	8.79	11.97	8.54	0.108	0.092	0.128	0.101
Precuneous cortex	7439	14.43	12.53	16.58	12.05	0.156	0.134	0.181	0.158
Cuneal cortex	1605	2.06	1.52	2.80	1.40	0.021	0.015	0.028	0.015
Frontal orbital cortex	4058	25.06	22.81	27.46	21.40	0.289	0.259	0.321	0.298
Parahippocampal gyrus (A)	3225	28.56	26.20	30.82	23.63	0.336	0.304	0.368	0.332
Parahippocampal gyrus (P)	2014	7.72	6.41	9.17	7.02	0.080	0.066	0.096	0.089
Lingual gyrus	5387	14.04	12.35	16.01	11.35	0.151	0.132	0.174	0.156
Temporal fusiform cortex (A)	830	1.40	1.00	1.87	1.05	0.014	0.010	0.019	0.013
Temporal fusiform cortex (P)	2728	5.45	4.53	6.57	4.91	0.056	0.046	0.068	0.055
Temporal occipital fusiform cortex	2451	17.51	15.48	19.63	15.56	0.193	0.168	0.219	0.212
Occipital fusiform gyrus	3576	14.36	12.58	16.30	11.81	0.155	0.134	0.178	0.146
Frontal operculum cortex	1059	6.91	5.84	8.03	5.85	0.072	0.060	0.084	0.065
Central operculum cortex	2568	7.54	6.46	8.83	6.32	0.078	0.067	0.092	0.070
Parietal operculum cortex	1683	3.35	2.55	4.27	3.39	0.034	0.026	0.044	0.039
Planum polare	1210	4.27	3.54	5.10	3.63	0.044	0.036	0.052	0.041
Heschl's gyrus	786	2.90	2.21	3.72	3.51	0.029	0.022	0.038	0.039
Planum temporale	1434	4.53	3.68	5.39	4.44	0.046	0.037	0.055	0.049
Supracalcarine cortex	424	0.78	0.55	1.08	0.94	0.008	0.005	0.011	0.009
Occipital pole	4081	15.36	13.27	17.56	12.40	0.167	0.142	0.193	0.175

Table 3: Meta-analysis results for executive control. Table presents posterior summaries for the probability of at least one activation in a given ROI (rounded to 2 decimal points), as well as the expected number points (rounded to 3 decimal points). The empirical quantities are also presented. Note the abbreviations: (L) for left, (R) for right, (A) for anterior, (P) for posterior, (S) for superior and (I) for inferior. All results are based on 1,000 HMC draws which are obtained after thinning the initial run every 10 iterations.

		$\% \, \mathbb{P}\left(\mathbf{N_X}\left(\mathbf{B}\right) \geq \mathbf{1}\right)$				$\int_{\mathbf{B}} \lambda\left(\xi\right) \mathbf{d}\xi$			
ROI	**VOL.**	$\mathbf{p_{0.5}}$	$\mathbf{p_{0.025}}$	$\mathbf{p_{0.975}}$	**DATA**	$\mathbf{p_{0.5}}$	$\mathbf{p_{0.025}}$	$\mathbf{p_{0.975}}$	**DATA**
Lateral ventrical (L)	1289	3.48	2.57	4.81	2.96	0.035	0.026	0.049	0.030
Lateral ventrical (R)	1019	2.65	1.90	3.54	2.37	0.027	0.019	0.036	0.024
Thalamus (L)	1591	14.65	11.84	18.26	13.02	0.158	0.126	0.202	0.154
Thalamus (R)	1398	11.56	8.89	14.41	10.95	0.123	0.093	0.156	0.127
Caudate (L)	572	3.39	2.36	4.73	3.25	0.034	0.024	0.048	0.033
Caudate (R)	515	2.60	1.73	3.65	2.37	0.026	0.017	0.037	0.024
Putamen (R)	800	4.73	3.36	6.44	5.62	0.048	0.034	0.067	0.059
Putamen (L)	923	5.69	4.21	7.54	5.33	0.059	0.043	0.078	0.053
Pallidum (L)	312	3.37	2.29	4.85	3.25	0.034	0.023	0.050	0.033
Pallidum (R)	266	2.48	1.56	3.71	2.66	0.025	0.016	0.038	0.027
Brain stem	7413	8.18	5.90	10.73	7.10	0.085	0.061	0.113	0.080
Hippocampus (L)	921	2.36	1.54	3.52	2.07	0.024	0.016	0.036	0.024
Hippocampus (R)	772	1.26	0.75	2.05	1.48	0.013	0.008	0.021	0.018
Amygdala (L)	390	0.84	0.43	1.48	0.59	0.008	0.004	0.015	0.006
Amygdala (R)	399	0.40	0.18	0.81	0.30	0.004	0.002	0.008	0.003
Accumbens (L)	111	0.52	0.27	1.00	1.48	0.005	0.003	0.010	0.015
Accumbens (R)	86	0.40	0.20	0.79	0.00	0.004	0.002	0.008	0.000
Frontal pole	15188	55.22	51.03	59.15	45.27	0.803	0.714	0.895	0.796
Insular cortex	3613	32.71	28.90	36.61	30.47	0.396	0.341	0.456	0.450

Superior frontal gyrus	7295	30.85	27.48	34.54	25.15	0.369	0.321	0.424	0.317
Middle frontal gyrus	6722	61.45	57.49	64.89	51.48	0.953	0.855	1.047	0.982
Inferior frontal gyrus, pars triangularis	1615	12.97	10.79	15.66	12.72	0.139	0.114	0.170	0.151
Inferior frontal gyrus, pars opercularis	1877	29.33	25.73	33.49	26.63	0.347	0.297	0.408	0.337
Precentral gyrus	11055	68.34	65.09	71.40	55.33	1.150	1.052	1.252	1.172
Temporal pole	3341	6.21	4.56	8.37	4.14	0.064	0.047	0.087	0.056
Superior temporal gyrus (A)	622	1.11	0.68	1.84	0.59	0.011	0.007	0.019	0.006
Superior temporal gyrus (P)	2093	6.76	5.06	8.72	5.62	0.070	0.052	0.091	0.074
Middle temporal gyrus (A)	762	0.86	0.44	1.61	0.59	0.009	0.004	0.016	0.009
Middle temporal gyrus (P)	2794	7.04	5.29	9.04	6.51	0.073	0.054	0.095	0.077
Middle temporal gyrus, temporooccipital part	2470	10.65	8.11	13.47	9.17	0.113	0.085	0.145	0.109
Inferior temporal gyrus (A)	789	0.89	0.43	1.61	1.48	0.009	0.004	0.016	0.015
Inferior temporal gyrus (P)	2687	4.38	3.09	6.18	2.96	0.045	0.031	0.064	0.041
Inferior temporal gyrus, temporoocipital part	1883	11.95	9.60	14.69	11.24	0.127	0.101	0.159	0.139
Postcentral gyrus	7553	19.37	16.25	22.68	13.91	0.215	0.177	0.257	0.210
Superior parietal lobule	3626	29.07	25.56	32.94	28.99	0.343	0.295	0.400	0.358
Supramarginal gyrus, (A)	1951	10.43	8.25	12.98	9.76	0.110	0.086	0.139	0.112
Supramarginal gyrus, (P)	3017	26.64	23.05	30.39	26.92	0.310	0.262	0.362	0.346
Angular gyrus	2944	26.03	22.79	29.84	22.49	0.301	0.259	0.354	0.290
Lateral occipital cortex (S)	9535	68.22	64.60	71.56	54.73	1.146	1.039	1.257	1.145
Lateral occipital cortex (I)	5004	25.77	22.04	29.63	20.41	0.298	0.249	0.351	0.314
Intracalcarine cortex	2211	8.35	6.41	10.62	8.88	0.087	0.066	0.112	0.101
Frontal medial cortex	1473	1.13	0.61	1.89	1.18	0.011	0.006	0.019	0.012
Juxtapositional lobule cortex	2193	31.78	27.80	36.38	30.18	0.382	0.326	0.452	0.391
Subcallosal cortex	2150	2.55	1.59	3.99	1.48	0.026	0.016	0.041	0.021

Paracingulate gyrus	4095	38.76	34.81	42.96	41.12	0.490	0.428	0.561	0.524
Cingulate gyrus (A)	4144	19.72	16.63	23.34	18.64	0.220	0.182	0.266	0.240
Cingulate gyrus (P)	4668	10.78	8.30	13.93	10.65	0.114	0.087	0.150	0.139
Precuneous cortex	7439	27.51	23.96	31.35	21.89	0.322	0.274	0.376	0.308
Cuneal cortex	1605	6.87	5.19	8.92	6.21	0.071	0.053	0.093	0.065
Frontal orbital cortex	4058	20.06	16.85	23.39	15.98	0.224	0.185	0.266	0.222
Parahippocampal gyrus (A)	3225	3.74	2.51	5.36	2.66	0.038	0.025	0.055	0.041
Parahippocampal gyrus (P)	2014	4.22	3.02	5.89	3.55	0.043	0.031	0.061	0.041
Lingual gyrus	5387	15.50	12.74	18.71	13.31	0.168	0.136	0.207	0.160
Temporal fusiform cortex (A)	830	0.87	0.48	1.60	0.89	0.009	0.005	0.016	0.012
Temporal fusiform cortex (P)	2728	4.55	3.27	6.09	3.25	0.047	0.033	0.063	0.036
Temporal occipital fusiform cortex	2451	13.68	11.08	16.69	10.95	0.147	0.117	0.183	0.136
Occipital fusiform gyrus	3576	20.13	17.26	23.72	15.98	0.225	0.190	0.271	0.225
Frontal operculum cortex	1059	15.83	13.44	18.52	12.72	0.172	0.144	0.205	0.148
Central operculum cortex	2568	11.05	8.91	13.37	10.36	0.117	0.093	0.143	0.115
Parietal operculum cortex	1683	4.73	3.44	6.51	3.55	0.048	0.035	0.067	0.047
Planum polare	1210	2.45	1.71	3.42	1.78	0.025	0.017	0.035	0.021
Heschl's gyrus	786	2.22	1.43	3.29	2.07	0.022	0.014	0.033	0.021
Planum temporale	1434	4.37	3.13	5.77	4.73	0.045	0.032	0.059	0.050
Supracalcarine cortex	424	1.20	0.82	1.77	1.48	0.012	0.008	0.018	0.015
Occipital pole	4081	17.34	14.09	20.92	13.31	0.190	0.152	0.235	0.186

E Model assessment

We use posterior predictive checks (Gelman *et al.*, 1996) to assess how well our model fits the data. First, we generate 1,000 realisations $\mathbf{y}^*_{kt}$ from the posterior predictive distributions $\mathbf{Y}^*_{kt} \sim \mathcal{PP}\left(\mathcal{B}, \boldsymbol{\lambda}^{kt}\right)$, where $k = 1$ for, $k = 2$ for executive control, and $\boldsymbol{\lambda}^{kt}$ is the t-th draw from the posterior intensity of type k. We then use these observations to construct several summary statistics which we compare with the values observed in the real data. Both first and second order properties of the point patterns are considered.

For first order properties we study $N_{\mathbf{Y}^*_{kt}}(B)$, the total number of points in some region B of the brain. For the entire brain, the posterior predictive samples yield 95% credible intervals of $[2, 13]$ for emotion and $[6, 20]$ for executively control. These cover 73% (623/855) and 63% (213/338) studies of the sample respectively. We further perform the same check for the 65 ROIs used in the full brain analysis of Section D. For each ROI and type, we record the percentage of studies that are contained in the 95% credible intervals of the posterior predictive distribution. Results are shown in Figure 16. We also find that from the 1193 studies of the sample, 1148 (96%) have at least 90% of their ROI counts covered by the posterior predictive intervals.

Our second test is based on the L-function (Møller and Waagepetersen, 2004, 2007) which is a summary statistic for the second order properties of a point pattern. For our model, given an observation $\mathbf{y}^*_{kt}$ and an intensity estimate $\boldsymbol{\lambda}^{kt}$ the L-fucntion at distance d can be estimated as:

$$L\left(d \mid \mathbf{y}^*_{kt}, \lambda^{kt}\right) = \left[\frac{3}{4\pi|\mathcal{B}|} \sum_{y_1, y_2 \in \mathbf{y}^*_{kt}} \frac{\mathbf{1}_{\{||y_1, y_2|| \leq d\}}}{\boldsymbol{\lambda}^{kt}_{\boldsymbol{\nu}_{y_1}}, \boldsymbol{\lambda}^{kt}_{\boldsymbol{\nu}_{y_2}}}\right]^{\frac{1}{3}}. \tag{28}$$

Therefore, as suggested by Illian *et al.* (2009), one can consider the differences $\Delta_{ik}(d) = L\big(d \mid \mathbf{x}_{ik}, \boldsymbol{\lambda}^{kt}\big) - L\big(d \mid \mathbf{y}^*_{kt}, \boldsymbol{\lambda}^{kt}\big)$ for all $t = 1, \ldots, 1,000$, where $\mathbf{x}_{ik}$ is the i-th observation of type k. If the 95% credible interval of Δ_{ik} for some study i does not contain zero then the model does a poor job explaining the data of this study. We investigate the differences $\Delta_{ik}(r)$ using values of r in a grid ranging from 0mm to 200mm, with spacing 2mm. Figure 17 summarises the results. In the top panel we see that the median (over studies) upper and lower bounds contain zero for both emotion and executive control studies. For each r, we calculate the proportion of studies that contain zero in the posterior predictive folders of Δ_{ik} and show results in the bottom panel of Figure 17. We find that for $r \leq 40mm$ over 85% of the studies of both types contain zero within the credible interval. For larger distances, the proportion drops but remains well above 60% suggesting no major issues with the fit of our model.

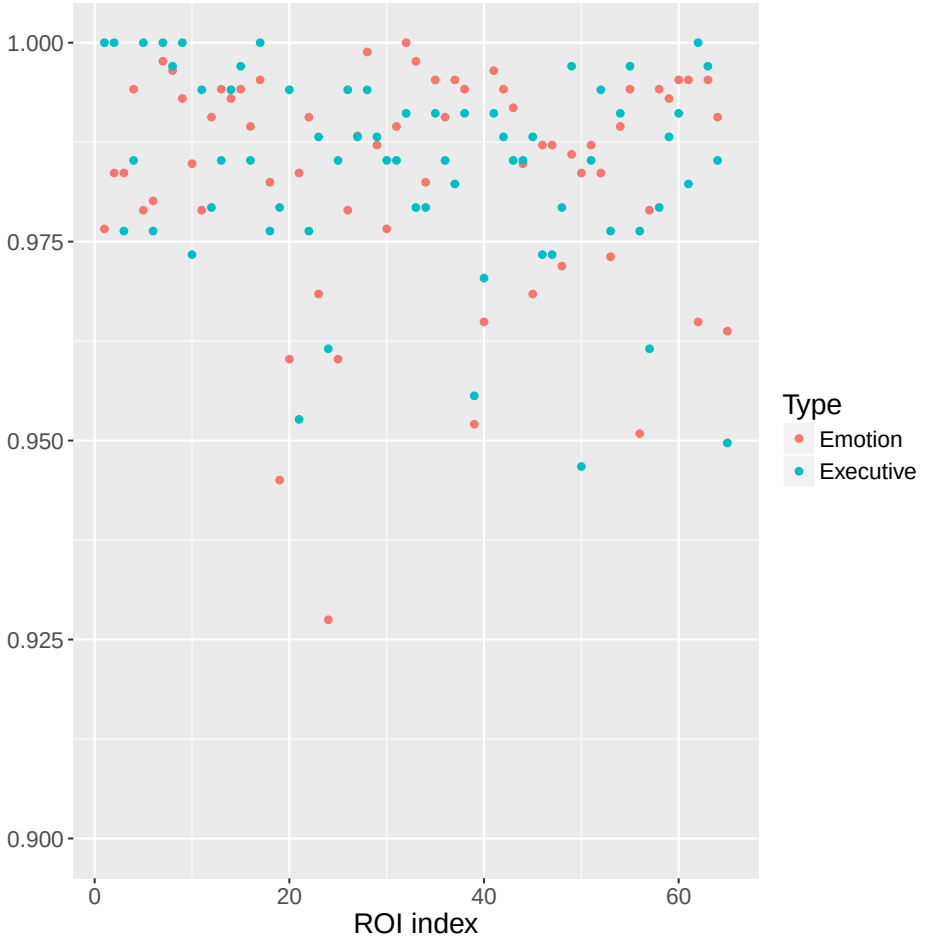

Figure 16: Results of the first posterior predictive check for first order propoerties. The x-axis is the ROI index in the same order presented in the full brain analysis of Section D. The y-axis represents the percentage of our meta-analysis studies that have counts which fall within the 95% credible interval, as obtained from the 1,000 posterior predictive samples. Emotion studies are shown in red whereas executive control studies are shown in blue. The posterior predictive samples were generated from 1,000 intensity draws obtained from the HMC.

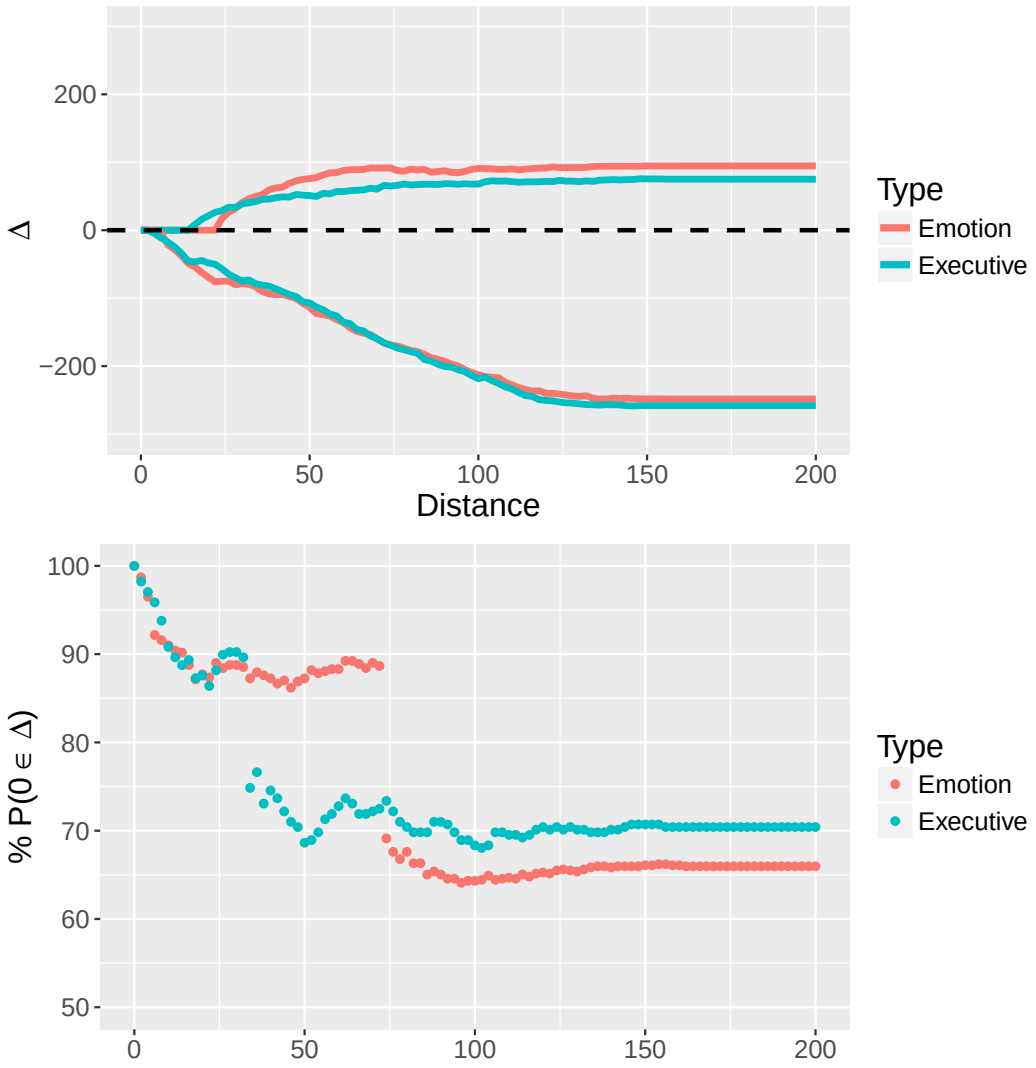

Figure 17: Results of the second posterior predictive check for second order propoerties. The upper panel shows the median over studies upper and lower credible intervals for Δ_{ik}. The bottom panel shows the proportion of studies which have credible intervals for Δ_{ik} that contain zero. Emotion studies are shown in red whereas executive control studies are shown in blue. The posterior predictive samples were generated from 1,000 intensity draws obtained from the HMC.